Privacy-Preserving and Accountable Multi-agent Learning

Anudit Nagar
Bennett University, India
anudit@bennett.edu.in

Cuong Tran
Syracuse University, New York
cutran@syr.edu

Ferdinando Fioretto
Syracuse University, New York
ffiorett@syr.edu

*A **robust and practical** framework for **privacy-preserving** and **accountable** multi-agent learning under a non-IID setting, a varying number of agents, and **strict privacy constraints**.*

INTRODUCTION

Distributed multi-agent learning enables agents to cooperatively train a learning model without requiring them to share their dataset.

This paper uses a decentralized computational environment that enables to train a differentially private multi-agent learning model, guaranteeing both privacy and trust. The resulting framework, called **Privacy-preserving and Accountable Distributed Learning (PA-DL)** relies on the **Ethereum blockchain**, that combines an immutable data storage with a Turing-complete computational environment.

The privacy requirement is enforced by ensuring that the learned model is differentially private. PA-DL uses a clipping approach on the model parameters and the privacy analysis relies on composition methods and the moment accountant for the **Sampled Gaussian Mechanism.**

Accountability is achieved by running the computation on the immutable blockchain combined with a decentralized procedure that validates the genuineness of the agent contributions to the model.

THE PA-DL FRAMEWORK

Private and Accountable Distributed Learning (PA-DL) is a fully distributed learning framework that ensures privacy and accountability while keeping the network bandwidth low. The framework is schematically shown below,

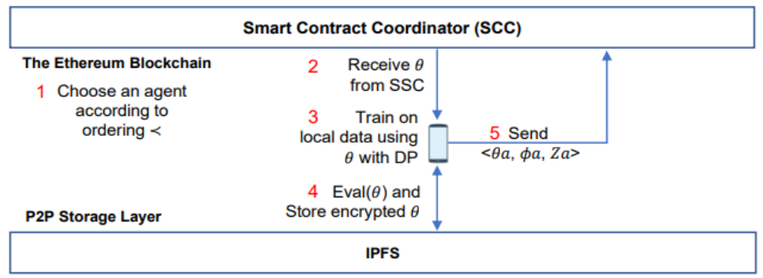

Flow diagram of the PA-DL Framework

- **A Smart Contract Coordinator (SCC):** It is a program executed on the blockchain that orchestrates the interaction among PA-DL agents to ensure the correct data exchange aimed at training a global model.
- **A provably private PA-DL agent training procedure:** At each round, the invoked agents use the parameters obtained by the SCC to train a model over their dataset. Each training step is ensured to guarantee (ϵ, δ)-differential privacy.
- **Accountability:** Prior being able to submit a model update, a PA-DL agent is required to invoke a verification step that ensures its trustworthiness.

EXPERIMENTAL RESULTS

Accuracy and Scalability

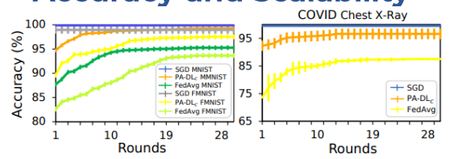

Agents	10	100	1000
PA-DL	0.102	1.020	10.20
FedAvg	0.600	6.000	60.00

Method	10^3 Tasks	10^6 Tasks
createTask	$2.76 * 10^{-3}$ USD	2.76 USD
startNextRound	$7.23 * 10^{-4}$ USD	0.723 USD

Algorithms' accuracy per round on MNIST & FMNIST (left), COVID-19 X-Ray(Right) for K = 1000 agents.

Network bandwidth in GB and Cost in USD (right) for execution of a method in SCC for N parallel tasks, required to complete 1 round on MNIST data

Privacy/Accuracy Trade-off

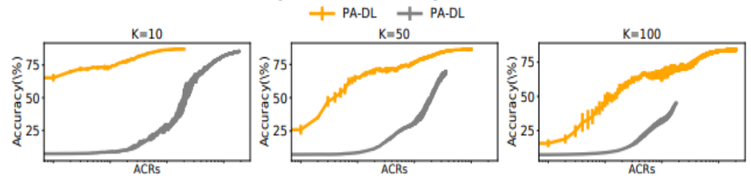

Accuracy on MNIST for K =10 (top), 50 (middle), and 100 (bottom) agents. The final privacy losses for the model with K =10, 50, and 100, respectively, are 0.5, 1.1 and 1.6.

Under a very tight privacy constraint $\epsilon \approx 1$, PA-DL consistently achieves significantly more accurate to those produced by DP-FedAvg. The experiment also analyzed biased datasets (not reported due to space constraints) and observed that the results follow the same trends as those outlined above.

Scan to **download** the **full paper.**