Distributed Learning Without Distress: Privacy-Preserving Empirical Risk Minimization

Bargav Jayaraman[1]; Lingxiao Wang[2]; David Evans[1]; Quanquan Gu[2]
[1]University of Virginia, [2]University of California Los Angeles

Abstract

Distributed learning allows a group of independent data owners to collaboratively learn a model over their data sets without exposing their private data. We present a distributed learning approach that combines differential privacy with secure multi-party computation.

We explore two popular methods of differential privacy, *output perturbation* and *gradient perturbation*, and advance the state-of-the-art for both methods in the distributed learning setting. In our output perturbation method, the parties combine local models within a secure computation and then add the required differential privacy noise before revealing the model. In our gradient perturbation method, the parties aggregate their local gradients within a secure computation, adding sufficient noise to ensure privacy before the gradient updates are revealed at each iteration. For both methods, we show that the noise can be reduced in the multi-party setting by adding the noise inside the secure computation after aggregation, asymptotically improving upon the best previous results.

Privacy-Preserving ERM in Single Party Setting

Three differential privacy mechanisms, all achieving the same lower bound on noise, which is inversely proportional to data set size.

$D = (X, y)$

Dataset Size: n

Model Training with Gradient Descent

$J(\theta) = \frac{1}{n}\sum_{i=1}^{n} \ell(\theta, X_i, y_i) + \lambda N(\theta) + \beta\left\{ \propto \frac{1}{n} \right\}$ Objective Perturbation [2]

for t in T :

$\theta \leftarrow \theta - \eta\left(\nabla J(\theta) + \beta\left\{ \propto \frac{1}{n} \right\}\right)$ Gradient Perturbation [1]

return $\theta + \beta\left\{ \propto \frac{1}{n} \right\}$ Output Perturbation [2]

θ

Privacy-Preserving ERM in Multi-Party Setting

In multi-party setting, the parties wish to collaboratively learn a machine learning model over their private data sets without revealing them to each other. This can be achieved by combining multi-party computation and differential privacy. However, in the multi-party setting, Pathak et al. [3] note that the bottleneck for differential privacy comes from the party which has the smallest data set.

Several approaches exist which can improve upon this bound by either combining private local models or generating distributed noise to be added to the final model. These approaches still do not achieve the optimal lower bound on the differential privacy noise that is possible in single party setting. Our proposed approaches achieve near optimal lower bound on noise in multi-party setting.

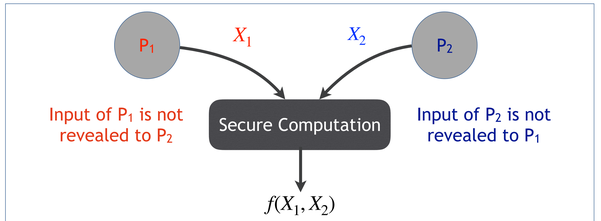

Generic Multi-Party Computation Protocol

Noise Comparison with Baseline Methods

Methods	Analytical Bound	Generated Noise*
Pathak [3]	Laplace$\left(\frac{2G}{n_{(1)}\lambda\epsilon}\right)$	1150 X 10^{-3}
Local Output Perturbation [2]	Laplace$\left(\frac{2G}{\sqrt{m}n_{(1)}\lambda\epsilon}\right)$	112 X 10^{-3}
Local Objective Perturbation [2]	Laplace$\left(\frac{2G}{n_{(1)}\epsilon'}\right)$	11.6 X 10^{-3}
Local Gradient Perturbation [5]	Gaussian$\left(\frac{\sqrt{2TG}}{\sqrt{m}n_{(1)}\epsilon}\right)$	5.63 X 10^{-3}
MPC Output Perturbation	Laplace$\left(\frac{2G}{mn_{(1)}\lambda\epsilon}\right)$	12.2 X 10^{-3}
MPC Gradient Perturbation	Gaussian$\left(\frac{\sqrt{2TG}}{mn_{(1)}\epsilon'}\right)$	0.572 X 10^{-3}

* $m = 100$, $n_{(1)} = 500$, $\lambda = 0.01$, $\epsilon = 0.5$, $G = 1$, $T = 100$
*Standard Deviations over 1000 samples

Output Perturbation for Multi-Party Setting

Smallest Dataset Size: $n_{(1)}$

D_1 D_2 $\cdots$ D_m

Model Training | Model Training | Model Training

$\theta^{(1)}$ $\theta^{(2)}$ $\theta^{(m)}$

MPC $\quad \theta = \frac{1}{m}\sum_{i=1}^{m} \theta^{(i)} + \beta\left\{ \propto \frac{1}{m \cdot n_{(1)}} \right\}$

θ

Gradient Perturbation for Multi-Party Setting

Model Training with Gradient Descent

for t in T :

Smallest Dataset Size: $n_{(1)}$

D_1 D_2 $\cdots$ D_m

$\nabla J_1(\theta)$ $\nabla J_2(\theta)$ $\nabla J_m(\theta)$

$\theta \leftarrow \theta - \eta \cdot \left(\frac{1}{m}\sum_{i=1}^{m} \nabla J_i(\theta) + \beta\left\{ \propto \frac{1}{m \cdot n_{(1)}} \right\} \right)$ MPC

return θ

θ

Experimental Results

We implemented our methods on output perturbation (**MPC Out P**) and gradient perturbation (**MPC Grad P**) along with the differential private baselines of Pathak et al. [3] (**Pathak**), Rajkumar and Agarwal [4] (**Rajkumar and Agarwal**), the multi-party extensions of Chaudhuri et al. [2] for output perturbation (**Local Out P**) and objective perturbation (**Local Obj P**), and an improved implementation of Shokri and Shmatikov [5] (**Local Grad P**).

We experimented with KDDCup99 dataset for classification and KDDCup98 dataset for regression, and report the relative performance of all methods with respect to a non-private baseline. Our methods perform better than the other private baselines for the same hyper-parameter settings. It is to be noted that while all the methods consume privacy budget of 0.5 for the entire training process, the method of Rajkumar and Agarwal consumes 0.5 privacy budget per iteration of model training.

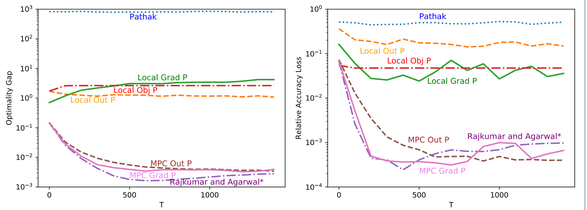

Performance Comparison on KDDCup99 Dataset for Classification (m = 1000)

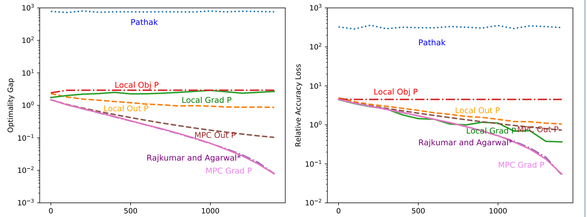

Performance Comparison on KDDCup98 Dataset for Regression (m = 1000)

Contact

Bargav Jayaraman
University of Virginia
Email: bargavjayaraman@gmail.com
Website: https://bargavjayaraman.github.io

Project Details

Website: https://oblivc.org/ppml/
Code: https://github.com/bargavj/distributedMachineLearning

References

1.Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar and Li Zhang. Deep learning with differential privacy. In *ACM Conference on Computer and Communications Security*, 2016.
2.Kamalika Chaudhuri, Claire Monteleoni and Anand D. Sarwate. Differentially private empirical risk minimization. In *Journal of Machine Learning Research*, 2011.
3.Manas Pathak, Shantanu Rane and Bhiksha Raj. Multiparty Differential Privacy via Aggregation of Locally Trained Classifiers. In *Advances in Neural Information Processing Systems*, 2010.
4.Arun Rajkumar and Shivani Agarwal. A differentially private stochastic gradient descent algorithm for multiparty classification. In *Artificial Intelligence and Statistics*, 2012.
5.Reza Shokri and Vitaly Shmatikov. Privacy-preserving deep learning. In *ACM Conference on Computer and Communications Security*, 2015.