2020「中技社科技獎學金」

2020 CTCI Foundation Science and Technology Scholarship

境外生生活助學金

Living Grant for International Graduate Students

Distributed Machine Learning with Differential Privacy

Yiwei Li, 3th Ph.D. student

Advisor: Prof. Chong-Yung Chi (祁忠勇)

Institute of Communications Engineering, National Tsing Hua University

Abstract

In distributed machine learning, the model is trained over distributed data sources through an interactive process of local computation and message passing. Such an iterative process could cause privacy concerns of data owners. We propose a novel framework based on the concept of differential privacy (DP), in which artificial noises are added to the parameters at the clients side before being uploaded to the server. In our work, first, we prove that the ADMM based distributed machine learning can satisfy DP under distinct protection levels by properly adapting the variances of artificial noises. Then we develop a theoretical convergence analysis on proposed model. Finally, we reveals that there exists a tradeoff between the convergence performance and privacy protection levels.

Research Focus

I. System Model

We consider the empirical risk minimization (ERM) problem in distributed machine learning, which can be formulated as follows:

$$\underset{\{\mathbf{w}_i\}}{\operatorname{argmin}} \sum_{i=1}^{n}\sum_{j=1}^{m_i}\frac{1}{m_i}\ell\left(\mathbf{x}_{i,j},\mathbf{y}_{i,j},\mathbf{w}_i\right)+\frac{\lambda}{n}R(\mathbf{w}_i) \tag{1a}$$

$$\text{s.t.} \quad w_i = w, i = 1,\dots,n \tag{1b}$$

where $\mathbf{w}_i$ is i-th local model, and $\mathbf{w}$ is the global one. $\ell(\cdot)$ is the loss function. $\mathbf{x}_{i,j}$ and $\mathbf{y}_{i,j}$ denote the j-th data feature and data label of i-th client, respectively. $R(\cdot)$ refers to the regularizer function and $\lambda > 0$ is the regularizer parameter. n and m_i respectively denote the number of clients and the number of local training samples. Then the augmented Lagrangian function associated with the problem (1) is

$$\mathcal{L}_\rho(\mathbf{w}_i,\mathbf{w},\boldsymbol{\gamma}_i)=\sum_{i=1}^{n}\mathcal{L}_{\rho,i}(\mathbf{w}_i,\mathbf{w},\boldsymbol{\gamma}_i) \tag{2}$$

$$\mathcal{L}_{\rho,i}(\mathbf{w}_i,\mathbf{w},\boldsymbol{\gamma}_i)=\sum_{j=1}^{m_i}\frac{1}{m_i}\ell\left(\mathbf{x}_{i,j},\mathbf{y}_{i,j},\mathbf{w}_i\right)+\frac{\lambda}{n}R(\mathbf{w}_i)-\gamma_i(\mathbf{w}_i-\mathbf{w})+\frac{\rho}{2}\|\mathbf{w}_i-\mathbf{w}\|^2 \tag{3}$$

where γ_i is the dual variable and $\rho > 0$ is the penalty parameter.

II. (ε, δ) - differential privacy

A randomized mechanism $\mathcal{M}$ is (ε,δ)-differential privacy if for any two neighboring data sets D and D' differing in only one record, and for any and for any output O of $\mathcal{M}$:

$$\mathbb{P}[\mathcal{M}(D)=O] \leqslant e^{\varepsilon}\mathbb{P}[\mathcal{M}(D')=O]+\delta \tag{4}$$

where ε is the privacy budget and δ is the probability that violate strict ε-differential privacy.

III. Distributed Learning Algorithm with (ε,δ)-differential privacy

In ADMM algorithm, the local updates interacting with server may be intercepted and this result in privacy leakage. Thus, we develop the difference privacy technique to protect local updates:

$$\text{Local update:}\quad \mathbf{w}_i^k \leftarrow \operatorname{argmin}_{\mathbf{w}_i}\mathcal{L}_{\rho,i}\left(\mathbf{w}_i,\mathbf{w}^{k-1},\boldsymbol{\gamma}_i^{k-1}\right)+\mathcal{N}_{d,p}\left(0,\sigma_{i,k}^2\mathbf{I}\right) \tag{5}$$

$$\text{Global update:}\quad \mathbf{w}^k \leftarrow \frac{1}{n}\sum_{i=1}^{n}\widetilde{\mathbf{w}}_i^k-\frac{1}{n}\sum_{i=1}^{n}\boldsymbol{\gamma}_i^{k-1}/\rho \tag{6}$$

where $\sigma_{i,k}^2$ is the power of norm gaussian noise in the k iteration on i-th client.

Research Results

I. Privacy Guarantee

Assume that the $\ell(\cdot)$ is L-smooth and the regularizer function $R(\cdot)$ is μ-strong convex, the ℓ_2 of $\ell(\cdot)$ is bounded by c_1. Our work prove that when the noise power $\sigma_{i,k}=2c_1\sqrt{2\ln(1.25/\delta)}/\left((\lambda/n+\rho)m_i\epsilon\right)$ can achieve (ε, δ)-differential privacy in each iteration. Moreover, we develop moments accountant to track the total privacy loss. The τ-th log moment of the total privacy loss in overall T iterations on i-th client is:

$$\alpha_i(\tau)=\sum_{k=1}^{T}\alpha_i^k(\tau)=\frac{T\tau(\tau+1)\epsilon^2}{4\ln(1.25/\delta)} \tag{7}$$

II. Convergence Analysis

We prove that our algorithm converges at a rate of $O(1/\sqrt{T})$ and by relaxing the privacy protection level, the convergence performance will be improved.

Note that this is our latest work and we will release more details later. You may follow my work via the link: https://ywei.jimdo.com/ .

Research Experience

- 2015/04 – 2017/08: Research Assistant, Quanzhou Institute of Equipment Manufacturing, Chinese Academy of Sciences.
- 2017/09 – Present: Ph.D. student, National Tsing Hua University.
- 2019/09 – Present: Visiting Ph.D. student, The Chinese University of Hong Kong (Shenzhen).

財團法人中技社 CTCI FOUNDATION