Not Just Cloud Privacy: Protecting Client Privacy in Teacher-Student Learning

Lichao Sun
Department of Computer Science
University of Illinois at Chicago
Chicago, IL 60607
lsun29@uic.edu

Ji Wang
College of Systems Engineering
National University of Defense Technology
Changsha, China 410073
wangji@nudt.edu.cn

Philip S. Yu
Department of Computer Science
University of Illinois at Chicago
Chicago, IL 60607
psyu@uic.edu

Lifang He
Department of Computer Science
Lehigh University
Bethlehem, PA 18015
lifanghescut@gmail.com

Abstract

Ensuring the privacy of sensitive data used to train modern machine learning models is of paramount importance in many areas of practice. One recent popular approach to study these concerns is using the differential privacy via a "teacher-student" model, wherein the teacher provides the student with useful, but noisy, information, hopefully allowing the student model to perform well on a given task. However, these studies only solve the privacy concerns of the teacher by assuming the student owns a public but unlabelled dataset. In real life, the student also has privacy concerns on its unlabelled data, so as to inquire about privacy protection on any data sent to the teacher. In this work, we re-design the privacy-preserving "teacher-student" model consisting of adopting both private arbitrary masking and local differential privacy, which protects the sensitive information of each student sample. However, the traditional training of teacher model is not robust on any perturbed data. We use the adversarial learning techniques to improve the robustness of the perturbed sample that supports returning good feedback without having all private information of each student sample. The experimental results demonstrate the effectiveness of our new privacy-preserving "teacher-student" model.

1 Local Private Data Protection via "Teacher-Student" Learning

Many attractive applications involve training models using highly sensitive data, to name a few, diagnosis of diseases with medical records or genetic sequences [2]. In order to protect the privacy of the training data, various privacy protection works have been proposed in the literature [12, 14, 20, 11]. The "teacher-student" learning framework with privacy constraints is of particular interest here, since it can provide a private student model without touching any sensitive data directly [7, 18, 15]. The original purpose of a teacher-student framework is to transfer the knowledge from the teacher model to help train a student to achieve similar performance with the teacher. To satisfy the privacy-preserving need, knowledge from the teacher model is carefully perturbed with random noise, before being passed to the student model. However, the current teacher-student frameworks (e.g. [15] and [16]) only take the privacy data concern of the teacher instead of the student.

To solve the privacy data concern of the student, we proposed a new framework name Local private Data Protection (LDP). The overview of LDP is presented in Fig. 1. LDP relies on the "teacher-
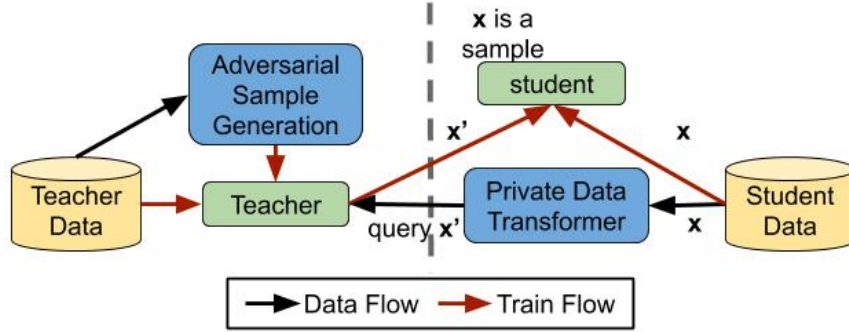
Figure 1: The overview of LDP.

student" model. The student needs to transform its private data with perturbed information for privacy protection locally. The teacher will use the adversarial sample generation to robust its performance to the perturbed sample from student data. In this way, the student will protect the privacy of its data, and the teacher could still return the correct label while facing the perturbed data. Before we introduce our protection mechanism in details, we first revise the definition of local differential privacy.

1.1 Local Differential Privacy

In this section, we revisit the definition of local differential privacy, which is a concept of privacy tailored to the privacy-preserving data analysis. It aims at providing provable privacy guarantee for each sensitive data sample, unlike general differential privacy is protecting the whole sensitive dataset [3]. Formally, the definition of ε-differential privacy is given as below:

Definition 1. *[5] A randomized mechanism $\mathcal{M}$ is ε-differential privacy, for any adjacent input x and x', and any output S of $\mathcal{M}$,*

$$\Pr[\mathcal{M}(x) = S] \leq e^{\varepsilon} \cdot \Pr[\mathcal{M}(x') = S]. \tag{1}$$

where the inputs x and x' are adjacent inputs when they differ by only one feature of each sample. The privacy guarantee of mechanism $\mathcal{M}$ is controlled by privacy budget [5], denoted as ε. A smaller value of ε indicates a stronger privacy guarantee. According to this definition, a differentially private algorithm can provide aggregate representations about a set of data items without leaking information of any data item.

A general method for approximating a deterministic function f with ε-differential privacy is to add noise calibrated to the global sensitivity of f, denoted as Δf, which is the maximal value of $\|f(x) - f(x')\|$ among any pair of x and x'. For instance, the Laplacian mechanism is defined by,

$$\mathcal{M}_f(x) = f(x) + Lap(\frac{\Delta f}{\varepsilon}), \tag{2}$$

where $Lap(\frac{\Delta f}{\varepsilon})$ is a random variable sampled from the Laplace distribution with scale $\frac{\Delta f}{\varepsilon}$.

The immunity to post-processing [6] also works on local differential privacy, which claims no algorithm can compromise the differentially private output and make it less differentially private.

1.2 Student: Private Data Transformer

To preserve the privacy of the student data, the local private data is perturbed by two advanced technologies: (1) private ranking masking can randomly hide the information of the private data; (2) local differential privacy injects additional perturbation which makes it hardly be estimated under local differential privacy guarantee.

Private Random Masking: It is one private random masking technique, which is an item-wise modification. First, we randomly generate a binary matrix I_n constituted of 0 and 1 with the same dimensions as input data x_s. The number of zeros in I_n is determined by $\lceil N \cdot \mu \rceil$, where $\lceil \cdot \rceil$ is the

---

Algorithm 1: Student: Private Data Transformer

Input: Each sensitive data $\boldsymbol{x}_s$.
Parameter :Random masking matrix I_n; Noise scale σ; Bound threshold B.

1   $\boldsymbol{x}_s \leftarrow \boldsymbol{x}_s \odot I_n$;

2   $\boldsymbol{x}_s \leftarrow \boldsymbol{x}_s / \max(1, \frac{\|\boldsymbol{x}_s\|_\infty}{B})$;

3   $\widetilde{\boldsymbol{x}}_s \leftarrow \boldsymbol{x}_s + Lap(B/\sigma\mathbf{I})$;

Output: Perturbed data $\widetilde{\boldsymbol{x}}_s$.

Algorithm 2: Teacher: Adversarial Sample Generation

Input: Teacher data $\boldsymbol{x}_r$.
Parameter :Noise scale σ; Bound threshold B; Controllers λ.

1   **foreach** $\boldsymbol{x}_r^{(i)}$ *in* $\{\boldsymbol{x}_r^{(1)}, ..., \boldsymbol{x}_r^{(N)}\}$ **do**

2       $\widetilde{\boldsymbol{x}}_r^{(i)} \leftarrow \boldsymbol{x}_r + Lap(B/\sigma\mathbf{I})$;

---

ceiling function, and μ is the random masking rate. The zeros are located in I_n conforming to the uniform distribution. Then, we use I_n item-wise multiplication on $\boldsymbol{x}_s$ for privacy protection, which can randomly hide the sensitive information out of the original data.

Local Differential Privacy: After private random masking, we adopt the local differential privacy for additional privacy protection. Hence, for each sensitive data $\boldsymbol{x}_s$, we clip the max value of the perturbed sample within fixed bounds to bound the sensitivity, *i.e.*, the output $\boldsymbol{x}_s$ is bounded as $\boldsymbol{x}_s / \max(1, \frac{\|\boldsymbol{x}_s\|_\infty}{B})$. It indicates that $\boldsymbol{x}_s$ is preserved when $\|\boldsymbol{x}_s\|_\infty \leq B$, whereas it is scaled down to B when $\|\boldsymbol{x}_s\|_\infty > B$. The bound threshold B is data-independent. The value of B is usually set as the median of the infinity norm of the original outputs over the whole data [1]. Then, we add the noise randomly sampled from the Laplace distribution into the bounded output $\widehat{\boldsymbol{x}}_s$ to protect the privacy.

1.3 Teacher: Adversarial Training via Noisy Sample Generation

The teacher model needs to use adversarial training to robust itself when it receives the perturbed sample from student. Without adversarial training, the teacher model can not give the correct label and feedbacks to the student, which hurts the performance while student wants to use the feedback information, such as student model training in Fig 1.

Adversarial Sample Generation: We generate the local differential private samples randomly from the teacher original data. In order to generate the same noisy scale perturbed samples, the teacher and student need to make an agreement on the same noise scale setting. Then, we can add the perturbed samples into the whole dataset for adversarial training as shown Algorithm 2.

Adversarial Training: The random noise added of student data is unknown to anyone except for the student itself. In this case, the teacher has no way to recover all information back of each student sample, due to the nature of randomness, which is obviously impractical to try on all possible random noisy perturbation. To enhance the robustness of the teacher model while facing the perturbed student sample, we train the teacher model under the worst situation via adversarial learning, which can be viewed as a min-max problem. We inject the worst perturbation $\boldsymbol{r}$ into the all data samples to maximize the model deviation from the original output, *i.e.*, the maximal loss $L(\boldsymbol{W}; \widetilde{\boldsymbol{x}}_r + \boldsymbol{r})$, while the deep neural network tries to minimize the deviation through training:

$$\min_{\boldsymbol{r};||\boldsymbol{r}||_2 \leq \eta} \max L(\boldsymbol{W}; \widetilde{\boldsymbol{x}}_r + \boldsymbol{r}), \tag{3}$$

where $\boldsymbol{r}$ is the maximum noise added to the whole data (*i.e.* including the Laplacian noise perturbed samples from adversarial sample generation), and η controls the scale of the noise.
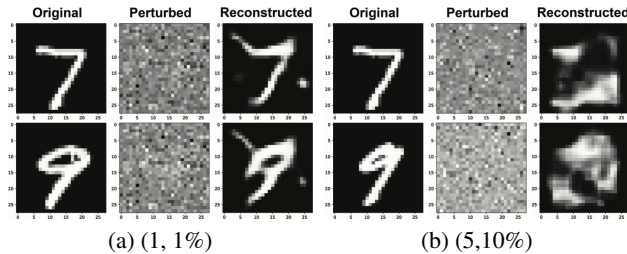
Figure 2: Visualization of noise and reconstruction.

2 Experimental Evaluation

In this section, image classification tasks are used as experimental examples to evaluate the effectiveness of LDP. We first examine the effect on three widely used image benchmark datasets, MNIST [10], SVHN [13], CIFAR-10 [8].

Meanwhile, we use Conv-Middle and Conv-Large for teacher training in LDP [9, 17]. For MNIST and SVHN, we use Conv-Middle as the teacher, and Conv-Large is used for CIFAR-10.

2.1 Student Evaluation: Privacy Protection

Here we visualize noise and reconstruction in Fig. 2 to demonstrate the effectiveness of private perturbation intuitively. While we set a very small perturbation (γ, μ) as $(1, 1\%)$, it already shows the privacy of the original data is partially protected. However, some existing works about image denoising and super-resolution reconstruction [19, 4] shows that it can reconstruct the original data from the perturbed samples. In this case, we train the model based on two perturbation strengths, and we found the small perturbation could be reconstructed, but already no longer as same as the original ones because of the benefits of the two private protection techniques. While we set the perturbation ratio as $(5, 10\%)$, the strength used in our experimental settings, the perturbed pictures with larger random noise is hardly reconstructed, which proves the privacy protection works well and the teacher hardly recover the original information of the student sample.

2.2 Teacher Evaluation: Performance Analysis

Here we verify the performance of the teacher model in LDP, we compare it with two baselines on three image datasets. Both BASE and LDP denotes the teacher model in Table **??**, *i.e.*, Conv-Middle for MNIST and SVHN, Conv-Large for CIFAR-10, but the BASE is trained without adversarial samples and adversarial learning. We test the BASE model in two situations: test without perturbation student data and test with perturbation student data. LDP is the complete framework proposed in this paper. λ is set as 0.2 for MNIST, SVHN, 0.5 for CIFAR-10.

Table 1: Performance Evaluation on Accuracy (%)

	Perturb	MNIST	SVHN	CIFAR-10
BASE	NO	99.21	95.42	86.37
BASE	YES	51.13	41.97	33.93
LDP	YES	98.16	90.02	79.52

Table 1 lists the result of every framework. BASE achieves the highest accuracy on no perturbed student data, but achieve the very bad performance on perturbed private student data. The testing accuracy of the BASE with perturbed data drops more than 40% average on three datasets compared with the BASE without perturbed data. These results indicate that the model trained on clean data only is not applicable to the prediction with perturbed data. In our case, the teacher is hard to maintain accurate feedbacks on the perturbed student data. Our proposed LDP adopts the adversarial learning with perturbed and clean data, then it can significantly mitigate the negative impact brought by the

private perturbation of the student data. As we can see, LDP can achieve almost close performance to the BASE model testing on non perturbed student data.

3 Conclusions

In this paper, we propose a new framework that first can protect the privacy of the student data in the "teacher-student" model. Then, We adopt adversarial data generation and adversarial learning techniques to enhance the robustness of the teacher on perturbed student data. The experimental results show the successful effeteness on both privacy protection of the student data and the performance of the teacher model evaluating the perturbed student data.

References

[1] Martín Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep Learning with Differential Privacy. In *Proceedings of the 23rd ACM SIGSAC Conference on Computer and Communications Security (CCS)*. 308–318.

[2] Babak Alipanahi, Andrew Delong, Matthew T Weirauch, and Brendan J Frey. 2015. Predicting the sequence specificities of DNA-and RNA-binding proteins by deep learning. *Nature biotechnology* 33, 8 (2015), 831.

[3] Amos Beimel, Hai Brenner, Shiva Prasad Kasiviswanathan, and Kobbi Nissim. 2014. Bounds on the sample complexity for private learning and private data release. *Machine Learning* 94, 3 (2014), 401–437.

[4] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. 2016. Image Super-Resolution Using Deep Convolutional Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38, 2 (2016), 295–307.

[5] Cynthia Dwork. 2011. *Differential Privacy*. Springer US, Boston, MA, 338–340.

[6] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science* 9, 3–4 (2014), 211–407.

[7] Jihun Hamm, Yingjun Cao, and Mikhail Belkin. 2016. Learning privately from multiparty data. In *International Conference on Machine Learning*. 555–563.

[8] A Krizhevsky and G Hinton. 2009. Learning Multiple Layers of Features from Tiny Images. (2009).

[9] Samuli Laine. 2017. Temporal Ensembling for Semi-Supervised Learning. In *5th International Conference on Learning Representations (ICLR)*.

[10] Yann LeCun, Léon Bottou, Yoshua Bengio, and Geoffrey Hinton. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.

[11] Xindi Ma, Jianfeng Ma, Hui Li, Qi Jiang, and Sheng Gao. 2018. PDLM: Privacy-Preserving Deep Learning Model on Cloud with Multiple Keys. *IEEE Transactions on Services Computing* (2018).

[12] Donald Michie, David J Spiegelhalter, CC Taylor, et al. 1994. Machine learning. *Neural and Statistical Classification* 13 (1994).

[13] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, and Bo Wu Andrew Y. Ng. 2011. Reading Digits in Natural Images with Unsupervised Feature Learning. In *Proceedings of the 24th International Conference on Neural Information Processing Systems (NIPS)*. 1–9.

[14] Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. 2007. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*. ACM, 75–84.

[15] Nicolas Papernot, Martin Abadi, Ulfar Erlingsson, Ian Goodfellow, and Kunal Talwar. 2017. Semi-supervised Knowledge Transfer for Deep Learning from Private Training Data. In *5th International Conference on Learning Representations (ICLR '17)*.

[16] Nicolas Papernot, Shuang Song, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Ulfar Erlingsson. 2018. Scalable Private Learning with PATE. In *6th International Conference on Learning Representations (ICLR '18)*.

[17] Sungrae Park, Jun-Keon Park, Su-Jin Shin, and Il-Chul Moon. 2017. Adversarial Dropout for Supervised and Semi-supervised Learning. *arXiv:1707.03631* (2017). https://doi.org/abs/1707.03631

[18] Manas Pathak, Shantanu Rane, and Bhiksha Raj. 2010. Multiparty differential privacy via aggregation of locally trained classifiers. In *Advances in Neural Information Processing Systems*. 1876–1884.

[19] Alec Radford, Luke Metz, and Soumith Chintala. 2015. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *arXiv:1511.06434* (2015). https://doi.org/abs/1511.06434

[20] Pouya Samangouei, Maya Kabkab, and Rama Chellappa. 2018. Defense-gan: Protecting classifiers against adversarial attacks using generative models. *arXiv preprint arXiv:1805.06605* (2018).