Personalized Differentially Private Federated Learning without Exposing Privacy Budgets

Anonymous Author(s)
Submission Id: 2886***

ABSTRACT

The meteoric rise of cross-silo Federated Learning (FL) is due to its ability to mitigate data breaches during collaborative training. To further provide rigorous privacy protection with consideration of the varying privacy requirements across different clients, a privacy-enhanced line of work on personalized differentially private federated learning (PDP-FL) has been proposed. However, the existing solution for PDP-FL [19] assumes the raw privacy requirements (i.e., privacy budgets) of all clients should be collected by the server, which are then *directly* utilized to improve the model utility via facilitating the privacy attitudes partitioning (i.e., partitioning all clients into multiple privacy groups). It is however non-realistic because the raw privacy budgets can be quite informative and sensitive.

In this work, our goal is to achieve PDP-FL without exposing clients' raw privacy budgets by indirectly partitioning the privacy attitudes solely based on clients' noisy model updates. The crux lies in the fact that the noisy updates could be influenced by two entangled factors of DP noises and non-IID clients' data, leaving it unknown whether it is possible to uncover privacy attitudes by disentangling the two affecting factors. To overcome the hurdle, we systematically investigate the unexplored question of *how to determine the conditions under which the model updates of clients can be dominated by the heterogeneous DP noises instead of non-IID data.* Then, we propose a simple yet effective strategy based on clustering the L2 norm of the noisy updates to indirectly estimate the privacy attitude partitions, which can be integrated into the vanilla PDP-FL to maintain the same performance. Experimental results demonstrate the effectiveness and feasibility of our privacy-budget-agnostic PDP-FL method.

CCS CONCEPTS

• **Security and privacy → Privacy protections**.

KEYWORDS

differential privacy, federated learning, personalization

ACM Reference Format:
Anonymous Author(s). 2023. Personalized Differentially Private Federated Learning without Exposing Privacy Budgets. In *Proceedings of (Conference acronym 'XX).* ACM, New York, NY, USA, 5 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Cross-silo Federated Learning (FL) [12, 21], which allows multiple clients to collaboratively train a global model without requiring access to clients' raw data, has been widely adopted both in academia and industry. Differential Privacy [5, 6] has been further integrated into FL, which gives rise to the DP-FL studies [2, 3, 7, 20, 24, 25] that seek to provide mathematically rigorous privacy protection at the desired level quantified by the privacy budget. DP-FL bears much resemblance to non-DP FL in training (e.g., by building on top of FedAvg [21]) but additionally incorporates local updates clipping and Gaussian noise injection [1, 4, 22, 28], whereby clients' local updates will be more strictly protected.

A more challenging yet practical problem is personalized differentially private federated learning (PDP-FL)[1], which takes the wide-ranging differences in individuals' privacy attitudes [11, 23, 26] into consideration and enables clients to pre-define their own privacy budgets (as opposed to shared an identical value specified by the server) [19]. Definition 1 formalizes this problem. One common way to achieve PDP in FL is to add different amounts of Gaussian noise to clients' submitted local updates, while directly aggregating the noisy and discordant local updates would inevitably lead to suboptimal model performance due to the biased estimation of the global parameters. To address these issues, Liu et al. [19] present the first promising attempt by developing a projection-based approach named projected federated averaging (PFA) for noise reduction [8, 30]. However, a major downside of PFA is that they treat clients' privacy budgets as publicly available knowledge and allow the server to utilize this information directly to identify the conservative/liberal clients at the initialization stage (see Line 5, Algorithm 1 in Section 2).

Definition 1 (Personalized Differential Privacy in Federated Learning [19])**.** Let the set of clients be $C = \{C_1, \ldots, C_M\}$, where each client $C_m \in C$ holds a local dataset $\mathcal{D}_m$. The federated learning satisfies $\{(\varepsilon_m, \delta_m)\}_{m\in[M]}$-personalized differential privacy, if each client satisfies $(\varepsilon_m, \delta_m)$-DP with respect to its local dataset.

We contend that assessing clients' privacy budgets is unrealistic and problematic. This is because the precise privacy budgets are also quite informative and sensitive for clients, and may act as a trigger for potential privacy attacks. Yet, we are not aware of any approach designed to discern the underlying privacy attitudes of clients based solely on their noisy model updates. Intuitively, the diversity of privacy budgets implies the varying magnitude of the perturbations added to the gradients, leading to a difference in the magnitude of clients' local updates due to the cumulative effects. On the other hand, such a difference also could be subject to the non-IID client data [15, 18]. An open question is *how to determine the*

[1] In our considered cross-silo setting, we use "personalized DP" to refer to customizing DP guarantees for each client rather than a specific user belonging to each client.

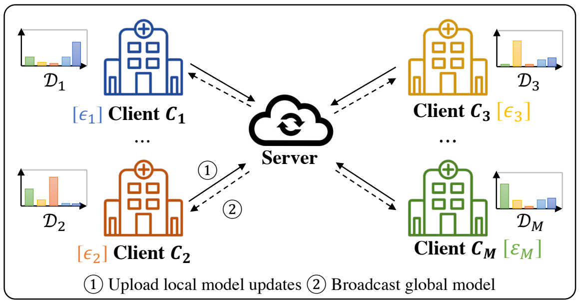

Figure 1: An illustration of the PDP-FL framework in which heterogeneous clients with non-IID data and personalized privacy budgets are collaboratively training a global model.

conditions under which the model updates of clients can be dominated by heterogeneous DP noises instead of non-IID data.

Contribution. In this paper, we aim to address the issue of indirectly estimating the privacy attitudes in the context of cross-silo FL for clients with non-IID data distributions and varying privacy budgets (ε). To summarize, our contributions are twofold.

(1) We discover through systematic empirical observations that the magnitude (i.e., L2-norms) of clients' local updates can serve as an effective indicator to facilitate indirect privacy attitudes partitioning. This novel insight propels the development of our clustering-based approach without requiring any prior knowledge about the real ε.

(2) We introduce a simple yet powerful approach for indirect privacy attitudes partitioning that suffices to leverage off-the-shelf clustering methods (e.g., Gaussian Mixture Models algorithm) to neglect the reliance on the raw privacy budgets in existing PDP-FL. To assess the effectiveness, we integrate it into the PFA framework and verify that our indirect privacy attitude partitioning approach can maintain the same model performance under the same experimental setup in the previous study[19].

2 PRELIMINARIES

Differential Privacy (DP). The definition of the classic (ϵ, δ)-DP is as follows, where the parameter ε is referred to as the *privacy budget* and the other parameter $\delta \geq 0$ captures the probability that the privacy is broken in the worst-case. A smaller value of ε corresponds to a higher level of privacy that can be achieved.

Definition 2 ((ε, δ)-Differential Privacy [6])**.** Let $\mathcal{D}$ be the space of all datasets and $D, D' \in \mathcal{D}$ is any pair of *adjacent datasets* where D' is obtained by deleting *any one* individual d from D, i.e., $D = D' \cup \{d\}$. A randomized mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{R}$ satisfies (ε, δ)-DP if for any subsets of outputs $S \subseteq \mathcal{R}$, it holds that

$$\Pr[\mathcal{M}(D) \in S] \leq e^{\varepsilon} \Pr[\mathcal{M}(D') \in S] + \delta.$$

Federated Averaging (FedAvg). FedAvg [21] is the most widely used algorithm for solving the federated optimization problem. In each communication round, a randomly sampled subset of clients run a certain number of Stochastic Gradient Descent (SGD) steps locally and independently, then the server averages the local updates and broadcasts a single global model to all clients. FedAvg

Algorithm 1: Projected Federated Averaging with Personalized Differential Privacy

input : Clients' privacy preferences $\{(\varepsilon_m, \delta)\}_{m\in[M]}$, number of communication rounds T, number of local steps τ
output: global model $\mathbf{x}_T$

Framework PDP-FL($\{(\varepsilon_m, \delta)\}_{m\in[M]}, T, \tau$):
  for *round* $t = 1, \ldots, T$ **do**
    $S_t \leftarrow$ (random subset of K clients)
    // Partition clients into "public" and "private"
    $S_t^{(pub)}, S_t^{(pri)} \leftarrow$
    (Before) Direct partition based on exposed privacy budgets $\{\varepsilon_m\}_{m\in S_t}$
    (After) Indirect partition based on clustering with L2-norms of the noisy local updates $\{\Delta\mathbf{x}_t^m\}_{m\in[K]}$
    foreach $m \in S_t$ **do** *in parallel*
      $\Delta\mathbf{x}_t^m \leftarrow$ DPSGD($t, \mathbf{x}_t, \tau$)
    $\Delta\bar{\mathbf{x}}_t \leftarrow$ PFA($\{(\varepsilon_m, \Delta\mathbf{x}_t^m)\}_{m\in[K]}, S_t^{(pub)}, S_t^{(pri)}$)
    $\mathbf{x}_{t+1} \leftarrow \mathbf{x}_t - \Delta\bar{\mathbf{x}}_t$
  return $\mathbf{x}_T$

Function PFA($\{(\varepsilon_m, \Delta\mathbf{x}^m)\}_{m\in[K]}, S^{(pub)}, S^{(pri)}$):
  // Compute the subspace from "public" updates
  $\mathbf{V}_k \leftarrow$ (The top-k eigenvectors of the second moment matrix computed from all $\Delta\mathbf{x}^m$ and $m \in \mathcal{S}^{(pub)}$)
  // Project "private" updates onto the subspace
  $\Delta\hat{\mathbf{x}}^{(pri)} \leftarrow \mathbf{V}_k\mathbf{V}_k^\top \sum_{m\in\mathcal{S}^{(pri)}} \omega_m \Delta\mathbf{x}^m$
  // Projected federated averaging
  $\mathcal{S} \leftarrow \mathcal{S}^{(pub)} + \mathcal{S}^{(pri)}$
  $\Delta\bar{\mathbf{x}} \leftarrow \frac{\sum_{m\in\mathcal{S}^{(pub)}} \varepsilon_m}{\sum_{m\in\mathcal{S}} \varepsilon_m} \cdot \Delta\bar{\mathbf{x}}^{(pub)} + \frac{\sum_{m\in\mathcal{S}^{(pri)}} \varepsilon_m}{\sum_{m\in\mathcal{S}} \varepsilon_m} \cdot \Delta\hat{\mathbf{x}}^{(pri)}$
  return $\Delta\bar{\mathbf{x}}$

by itself makes no special adjustments when encountering non-IID client data and therefore suffers from suboptimal performance in such circumstances [9, 16, 17].

Projected Federated Averaging (PFA). In PFA [19], all clients are divided into two types according to their precise privacy budgets (i.e., "private" clients with stricter privacy budgets and "public" clients with more relaxed privacy budgets) exposed to the server at the initialization stage; then the server extracts a reduced-dimensional subspace from the "public" model updates and projects the "private" model updates onto it. In this way, the heavy private perturbation of the "private" updates can be discarded, and the most useful information from all clients can be aggregated to improve the joint model utility. Pseudocode is given in Algorithm 1.

3 STUDY ON THE IMPLICATIONS OF NOISE LEVEL AND DATA DISTRIBUTION

In this section, we conduct a comprehensive empirical study to explore the characteristics of the local model updates obtained from heterogeneous clients whose local data distribution and privacy budgets differ from one another. This investigation aims to gain insights into the conditions under which the DP perturbations can significantly affect the client's local updates compared to the effect caused by non-IID data.

3.1 Experimental Setup

Datasets and Models. We consider two classic image classification tasks: the MNIST [14] digit recognition with a simple logistic regression model (MNIST-LogR) and the CIFAR10 [13] image classification with the same CNN architecture as in McMahan et al. [21]. We deploy them in the cross-silo FL setting with $M = 10$ clients.

Data Distributions. To examine the effects of data heterogeneity, we first establish the baseline using IID data and consider two partition strategies to simulate potential non-IID scenarios.

- **IID**: each client is assigned a uniform distribution over 10 classes.
- **NIID(2)**: also known as the *quantity-based label distribution skew* where each client owns data records of a fixed number (e.g., 2) of labels [21].
- **NIID-Dir(0.5)**: also known as the *distribution-based label imbalance* where a $p_{k,m} \sim Dir(\beta)$ proportion of records of class k are allocated to client m. Here $Dir(\beta)$ denotes a Dirichlet distribution [10] and the smaller the β is, the resulting partition is more unbalanced. We choose the same $\beta = 0.5$ as done in [29].

Varying privacy budgets. We explore a diverse range of privacy budgets ε to manifest the significant differences in privacy requirements among clients with varying privacy attitudes (e.g., $\varepsilon \approx 0.4, 3.0, 20$ for the MNIST-LogR experiments), and establish the baseline without any DP requirement.

Methods. For all experiments, we employ FedAvg [21] as the base FL algorithm. To ensure DP-FL, we incorporate minibatch DP-SGD [1] into clients' local training procedures, resulting in a modified version of FedAvg known as DP-FedAvg. In brief, DP-FedAvg introduces a certain amount of Gaussian noise to the clipped gradients during each local SGD iteration. It is worth noting that we do not employ the PFA algorithm since the privacy budgets of all clients are hidden from the server side, making the projection-based operations inapplicable in this case. Furthermore, we incorporate the full participation procedure to ensure all clients get continuous observations throughout the communication rounds.

Evaluation Metrics. In this section, we always report the average and standard deviation of the L2-norms of local updates across all clients along the training process. For the sake of readability, we use the abbreviations *avg./std. L2-norm* in the remaining sections.

Hyperparameters. Unless otherwise stated, we fix the local minibatch size $B = 64$, the local epochs $E = 1$, the total number of communication rounds $T = 20$, and the step size (or learning rate) $\eta = 0.01$ for all clients.

3.2 Evaluation Results

For every single plot in Fig. 2, we show how the avg./std. L2-norm evolves over communication rounds in the IID, NIID(2) and NIID-dir(0.5) settings respectively. Furthermore, we conduct a series of comparative experiments of FL with/without DP to analyze the isolated implications of varying levels of additive Gaussian noise on the values of avg./std. L2-norm. Here we assume that all clients have identical privacy budgets, which we refer to as homogeneous DP (HomoDP) in contrast to PDP. The intention behind this consideration is to explore the differences in the characteristics of the avg./std. L2-norms among clients with different privacy budgets.

The isolated effect of data distribution. From Fig. 2, we can observe two common trends from all plots: (1) both the avg. and the std. L2-norms in IID cases consistently exhibit lower values compared to all non-IID cases along the training process; (2) in the majority of cases, NIID-Dir(0.5) tends to produce avg. L2-norms

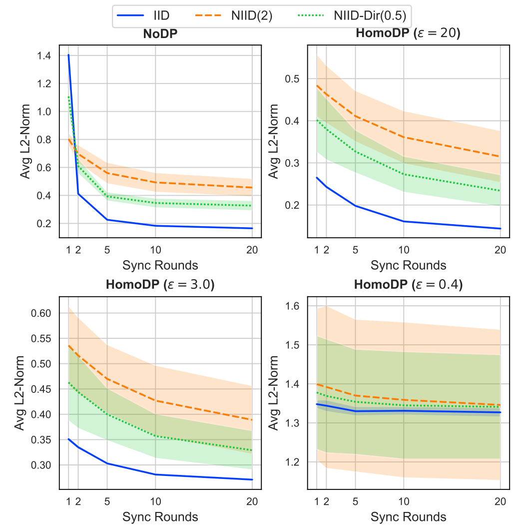

(a) Experimental results evaluated on MNIST-LogR

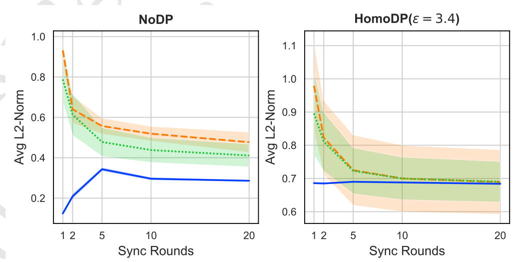

(b) Experimental results evaluated on CIFAR10-CNN

Figure 2: Effects of data distribution and varying privacy budgets on the average and standard deviation of the L2-norms (y-axis) of local updates across 10 clients over a maximum of 20 communication rounds (x-axis).

and std. L2-norms that are either smaller or comparable to those obtained with NIID(2).

The isolated effect of privacy budget. From Fig. 2 (a), it is clear that there exists a negative correlation between the value of privacy budget (ε) and the avg./std. L2-norms in two non-IID cases. Although the trend may not be readily apparent in the IID case, we note that the observation remains consistent. It makes sense since the discrepancies in privacy budgets imply variations in the scale of the random Gaussian distribution, resulting in different amounts of additive noise being introduced to the model updates during the local training procedure. Surprisingly, the results obtained from the cases with $\varepsilon = 3.0$ and $\varepsilon = 20$ show a considerable resemblance, indicating that both cases result in a similar degree of perturbation on the magnitude of clients' local updates, despite the latter case having a significantly larger privacy budget (in other words, an ε value of 3.0 may not be sufficiently small to provide a significant enhancement in privacy protection compared to the weak privacy setting of $\varepsilon = 20$). Given that similar trends have been observed in the CIFAR10-CNN experiments, we present only a partial set of results here due to the strict space limitations.

Table 1: Distribution of privacy preferences

Distribution	Parameters Setting
MixGauss1	Mixture of $\mathcal{N}_1(0.1, 0.01)$ and $\mathcal{N}_2(10.0, 0.1)$ with mixture weights 0.9 and 0.1
MixGauss2	Mixture of $\mathcal{N}_1(1.0, 0.1)$ and $\mathcal{N}_2(10.0, 0.1)$ with mixture weights 0.9 and 0.1
MixGauss3	Mixture of $\mathcal{N}_1(0.1, 0.01)$, $\mathcal{N}_2(1.0, 0.1)$ and $\mathcal{N}_3(10.0, 1.0)$ with mixture weights 0.5, 0.4 and 0.1

4 PDP-FL WITHOUT EXPOSING RAW PRIVACY BUDGETS

In this section, we introduce a privacy-budget-agnostic version of PFA that utilizes the L2-norms of noisy local updates. To evaluate the effectiveness of our approach as well as ensure a fair comparison, we reproduce the experiments using the same experimental setup as the previous study conducted by Liu et al [19].

4.1 Indirect Privacy Attitudes Partitioning

Key Insight. In our empirical study presented in the above section, we investigate the effects of non-IID data and varying privacy budgets on the local model updates of clients. The experimental results suggest that it is possible to indirectly partition the privacy attitudes of clients into groups by analyzing the L2-norms of their local noisy updates without requiring access to their raw privacy budgets, as long as (1) there exists a significant diversity in the privacy budgets across all clients; (2) the "private" (or conservative) clients opt for a ε that is small enough to ensure effective differentiation.

Proposed Approach. Equipped with the above key insight, now our focus shifts back to the PDP-FL setting where the additive Gaussian noises of the clients are drawn from different distributions determined by their privacy budget. The conservative clients with stricter privacy budgets require larger perturbation while the liberal clients with more relaxed privacy budgets submit more accurate model updates. This distinction in privacy budgets and the corresponding impact on the magnitude of perturbations just align with the two conditions revealed in the key insight from Section 3, which motivates us to develop the clustering-based approach for indirect privacy attitude estimation using L2-norms.

In more detail, we develop a simple yet powerful strategy based on the clients' noisy local updates and the Gaussian Mixture Models (GMMs) clustering algorithm, based on the intuition that clients who have similar privacy attitudes (privacy budgets) are expected to introduce Gaussian noises drawn from a similar random distribution. Then we can improve PFA by replacing the original *direct* client partition based on exposed privacy budgets with the *indirect* partition based on L2-norm clustering (as highlighted in Alg. 1).

4.2 Experimental Results

We first evaluate the utility of our proposed clustering approach by considering 3 potential multimodal distributions (a mixture of two or three different Gaussian distributions) as shown in Tab. 1 (see more details in [19]). Note that this assumption is supported by previous observations which have shown that a bimodal distribution is quite universal in a wide range of complex social systems [27].

Effects of the privacy preference distribution. In Fig. 3, we demonstrate the effectiveness of the L2-norm clustering approach evaluated on MNIST-LogR in NIID(2) setting with 10 clients in three privacy preferences distributions. In all plots, we utilize various markers to represent the predicted cluster index and the real privacy attitude of each client. Additionally, we use three different colors to indicate the resulting partitions. Experiment results show an obvious consistency between the L2-norms clustering and the ground truths (based on clients' real privacy budgets).

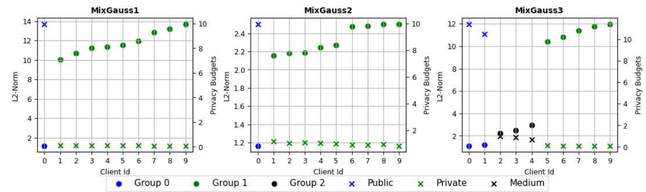

Figure 3: The consistency between the results of GMMs clustering based on L2-norm (left y-axis) and the ground truths based on the real privacy budgets (right y-axis) across 10 clients (x-axis) evaluated on MNIST-LogR in NIID(2) setting.

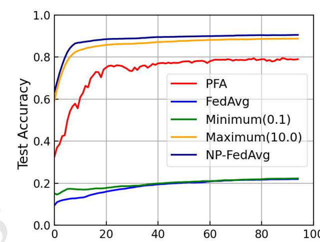

Figure 4: The test accuracy versus communication rounds evaluated on MNIST-LogR in non-IID data setting with privacy preferences distribution of MixGauss1.

Evaluation of the end-to-end PFA framework. In Fig. 4, we report the test accuracy versus communication rounds evaluated on MNIST-LogR in non-IID data setting with privacy preferences distribution of MixGauss1. Different from Liu et al. [19], we do not compare the weighted average (WeiAvg) and the communication-efficient version of PFA (PFA+) here since these two methods are dependent on the values of clients' privacy budgets, which is no longer available in our considered scenario. Just as we expected, the distinct utility advantages of PFA over the baseline methods FedAvg and Minimum remain due to the correct clustering results. Although it has worse accuracy than the non-private baseline (NP-FedAvg), PFA still reaches a reasonable level of model utility, while the FedAvg with PDP becomes ineffective. Just as we expected, the distinct utility advantages of PFA over the baseline methods FedAvg and Minimum remain due to the correct clustering results. Although it has worse accuracy than the non-private baseline (NP-FedAvg), PFA still reaches a reasonable level of model utility, while the FedAvg with PDP becomes ineffective.

5 CONCLUSION AND FUTURE WORK

In this work, we have proposed an effective method for indirect privacy attitude estimation based on L2-norm clustering in the PDP-FL setting. Additionally, we have integrated this clustering approach into the vanilla PFA framework to address potential privacy leakage issues arising from exposed privacy budgets. Some future directions include: (1) generalizing the clustering strategy to the more challenging cases where clients' privacy budgets are relatively uniform or more difficult to differentiate; (2) conducting extensive empirical evaluations on larger and more diverse datasets for deeper explorations into the effectiveness and scalability of our proposed approach.

REFERENCES

[1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *CCS*.

[2] Naman Agarwal, Peter Kairouz, and Ziyu Liu. 2021. The skellam mechanism for differentially private federated learning. In *NeurIPS*.

[3] Naman Agarwal, Ananda Theertha Suresh, Felix Yu, Sanjiv Kumar, and H. Brendan McMahan. 2018. CpSGD: Communication-Efficient and Differentially-Private Distributed SGD. In *NeurIPS*.

[4] Xiangyi Chen, Steven Z Wu, and Mingyi Hong. 2020. Understanding gradient clipping in private sgd: A geometric perspective. In *NeurIPS*.

[5] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *TCC*.

[6] Cynthia Dwork, Aaron Roth, et al. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science* 9, 3-4 (2014), 211–407.

[7] Robin C Geyer, Tassilo Klein, and Moin Nabi. 2017. Differentially private federated learning: A client level perspective. In *arXiv preprint arXiv:1712.07557*.

[8] Xin Gu, Gautam Kamath, and Zhiwei Steven Wu. 2023. Choosing public datasets for private machine learning via gradient subspace distance. In *arXiv preprint arXiv:2303.01256*.

[9] Filip Hanzely, Slavomír Hanzely, Samuel Horváth, and Peter Richtárik. 2020. Lower bounds and optimal algorithms for personalized federated learning. In *NeurIPS*.

[10] Jonathan Huang. 2005. Maximum likelihood estimation of Dirichlet distribution parameters. *CMU Technique report* 18 (2005).

[11] Zach Jorgensen, Ting Yu, and Graham Cormode. 2015. Conservative or liberal? Personalized differential privacy. In *ICDE*.

[12] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. 2021. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning* 14, 1–2 (2021), 1–210.

[13] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images.

[14] Yann LeCun, Corinna Cortes, and CJ Burges. 2010. MNIST handwritten digit database. *ATT Labs [Online]. Available: http://yann.lecun.com/exdb/mnist* 2.

[15] Qinbin Li, Yiqun Diao, Quan Chen, and Bingsheng He. 2022. Federated learning on non-iid data silos: An experimental study. In *ICDE*.

[16] Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. 2021. Ditto: Fair and robust federated learning through personalization. In *ICML*.

[17] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. 2020. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems* 2 (2020), 429–450.

[18] Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. 2019. On the convergence of fedavg on non-iid data. In *arXiv preprint arXiv:1907.02189*.

[19] Junxu Liu, Jian Lou, Li Xiong, Jinfei Liu, and Xiaofeng Meng. 2021. Projected federated averaging with heterogeneous differential privacy. In *VLDB*.

[20] Ken Liu, Shengyuan Hu, Steven Z Wu, and Virginia Smith. 2022. On privacy and personalization in cross-silo federated learning. In *NeurIPS*.

[21] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*.

[22] Milad Nasr, Reza Shokri, and Amir Houmansadr. 2019. Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. In *S&P*.

[23] Ben Niu, Yahong Chen, Boyang Wang, Zhibo Wang, Fenghua Li, and Jin Cao. 2021. AdaPDP: Adaptive personalized differential privacy. In *INFOCOM*.

[24] Maxence Noble, Aurélien Bellet, and Aymeric Dieuleveut. 2022. Differentially private federated learning on heterogeneous data. In *AISTATS*.

[25] Yifan Shi, Yingqi Liu, Kang Wei, Li Shen, Xueqian Wang, and Dacheng Tao. 2023. Make landscape flatter in differentially private federated learning. In *ICCV*.

[26] Zhibo Wang, Jiahui Hu, Ruizhao Lv, Jian Wei, Qian Wang, Dejun Yang, and Hairong Qi. 2018. Personalized privacy-preserving task allocation for mobile crowdsensing. *IEEE Transactions on Mobile Computing* 18, 6 (2018), 1330–1341.

[27] Ye Wu, Changsong Zhou, Jinghua Xiao, Jürgen Kurths, and Hans Joachim Schellnhuber. 2010. Evidence for a bimodal distribution in human communication. *PNAS* 107, 44 (2010), 18803–18808.

[28] Lei Yu, Ling Liu, Calton Pu, Mehmet Emre Gursoy, and Stacey Truex. 2019. Differentially private model publishing for deep learning. In *S&P*.

[29] Mikhail Yurochkin, Mayank Agarwal, Soumya Ghosh, Kristjan Greenewald, Nghia Hoang, and Yasaman Khazaeni. 2019. Bayesian nonparametric federated learning of neural networks. In *ICML*.

[30] Yingxue Zhou, Steven Wu, and Arindam Banerjee. 2021. Bypassing the Ambient Dimension: Private SGD with Gradient Subspace Identification. In *ICLR*.

Received 16 June 2023