Privacy and Personalization in Heterogeneous Federated Cross-Silo Learning Environments

Tanvi Shegaonkar
shegao_t1@denison.edu

Stacey Truex
truexs@denison.edu

Heterogeneous, Cross-Silo FL

Cross- Silo Setting

Consider the cross-silo setting in a federating learning environment

- Each hospital $i \ldots k$ represents a silo where each client may have sufficient data to train a sufficient local model without participating in FL
- Each record in a silo corresponds to a single user
- **Q:** When does participation in federated learning and/or personalization remain beneficial given data, trust, and privacy heterogeneity?

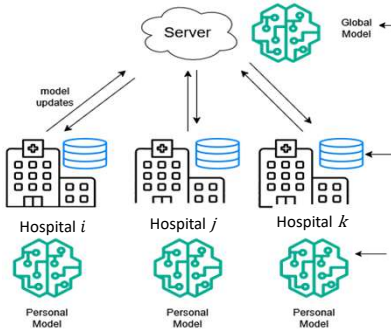

Global model is trained by aggregating model updates across silos

Queries to **each record** (x_i, y_i) are protected with differential privacy guarantee ε_i

Personal Models are trained with Ditto[1] The personal objective for each client k is $min_{v_k} h_k(v_k, w^*) := F_k(v_k) + \frac{\lambda}{2}||v_k - w^*||^2$

Sources of Heterogeneity

Data Heterogeneity : Clients may have data that is non-Independent and Identically Distributed (Non-IID) and suffer poor performance under a shared global model. We evaluate five systematically varied label skew distributions for MNIST, CIFAR-10 and Fashion MNIST:

- IID
- Non-IID-5 (50% of classes per client)
- Non-IID-2 (20% of classes per client)
- Dirichlet-0.1 ($\alpha = 0.1$)
- Dirichlet-0.5 ($\alpha = 0.5$)

E.g: Non-IID-2 label distribution

Definition: Record-level Differential Privacy (DP). For each data point (x_i, y_i) of client i, there is an assigned privacy parameter ε_i. This is enforced using DP-SGD [1] with the Gaussian mechanism[2] and privacy loss is tracked via Rényi DP[3]

Trust Model Heterogeneity:

- Non-private: No noise is added as server and clients are trusted
- External privacy: DP is applied to shared FL updates but not to local/personal training
- Total privacy: DP is applied to all training

Privacy Heterogeneity : A client's dataset exhibits privacy heterogeneity if $\exists j, k$ such that $\epsilon_j \neq \epsilon_k$.

- Certain labels can encode sensitive information, requiring stricter privacy guarantees
- **Data-dependent privacy**: privacy budgets are assigned from the 3-level distribution to data points in sorted class label order, with lower-indexed classes assigned tighter budgets.

$$\epsilon \sim \begin{cases} 0.1 & \text{with probability } 0.7 \\ 1.0 & \text{with probability } 0.2 \\ 5.0 & \text{with probability } 0.1 \end{cases}$$

Three-Levels Distribution [4]

Participation Incentives

Data Heterogeneity

FL Gain = Global Model Accuracy – Local Model Accuracy

Dataset	MNIST	Fashion-MNIST	CIFAR-10
Avg. FL Gain (Local Eval)	4.24± 0.93	9.38 ± .86	27.34± 2.89
Avg. FL Gain (Global Eval)	3.64 ± .96	9.37 ±.64	27.37±2.43

Table 1: Average accuracy gain from federated learning across datasets, evaluated on local and global data in an IID setting.

Dataset	Fashion-MNIST	CIFAR-10	Heart Disease
Avg. FL Gain (Local Eval)	-15.13 ± 3.70	-26.60 ± 26.94	-7.59 ± 13.41
Avg. FL Gain (Global Eval)	26.65 ± 2.31	12.36 ± 5.15	7.40 ± 9.80

Table 2: Avg FL accuracy gain under local/global evaluation on Non-IID-5 Fashion-MNIST and CIFAR-10 and Non-IID Heart Disease.

Local evaluation: accuracy over a client's local data

- clients have little incentive to participate in FL
- global model **does not provide good performance for every client** vs. local model in non-IID settings

Global evaluation: accuracy over all data across clients

- clients are **more strongly incentivized in non-IID** settings
- local datasets poorly approximate the population distribution.

Trust Model Heterogeneity

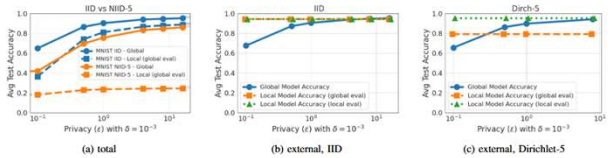

(a) total (b) external, IID (c) external, Dirichlet-5

Non-private:

- FL global model consistently outperform local models under global evaluation as shown in Tables 1 and 2

External privacy:

- IID setting: noise-free local training ⇒ **noise added to shared updates for ε-dp guarantee reduces the incentive to participate in FL** as local models outperform the global model
- Non-IID setting: data heterogeneity ⇒ **the incentive to participate in FL remains** as local models perform well locally but poorly globally. At low ε, degradation in global model accuracy exceeds the local model utility loss due to heterogeneity.

Total privacy:

- IID & Non-IID: Incentive to participate in FL remains as global models benefit from averaging of noisy updates [5]

Personalization & Privacy Heterogeneity

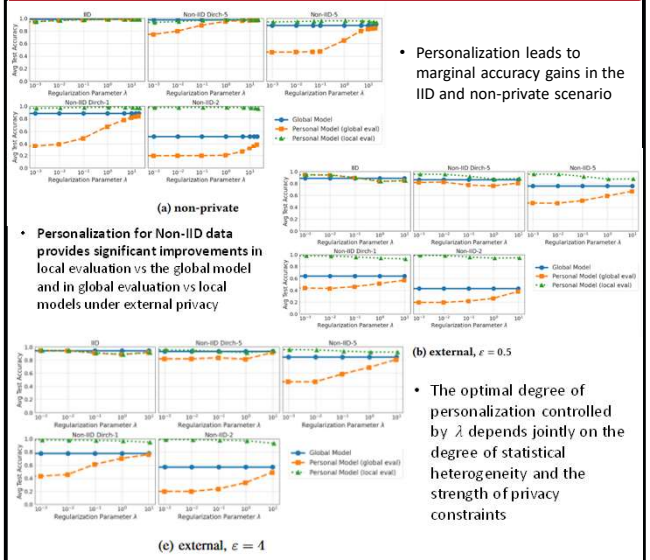

(a) non-private (b) external, $\varepsilon = 0.5$ (c) external, $\varepsilon = 4$

- Personalization leads to marginal accuracy gains in the IID and non-private scenario
- **Personalization for Non-IID data provides significant improvements in** local evaluation vs the global model and in global evaluation vs local models under external privacy
- The optimal degree of personalization controlled by λ depends jointly on the degree of statistical heterogeneity and the strength of privacy constraints

Data Dependent Privacy Heterogeneity

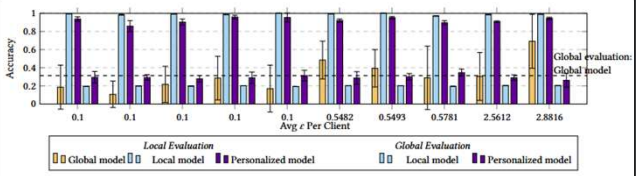

- Ordering clients by average ε reveals a **systematic performance disparity**: low-ε clients experience worse global performance than ε clients.
- **Personalized models achieve comparable globally evaluated accuracy** as the global model, **while mitigating disparities in performance between low and high ε clients** found in the global model.

Acknowledgements

Research conducted by the TLS lab at Denison University directed by Dr. Stacey Truex

[1] Martin 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 2016 ACM SIGSAC conference on computer and communications security. 308–318.
[2] Cynthia Dwork and Aaron Roth. 2014. The algorithmic foundations of differential privacy. Foundations and trends® in theoretical computer science 9, 3-4 (2014), 211–487.
[3] Ilya Mironov. 2017. Rényi differential privacy. In 2017 IEEE 30th computer security foundations symposium (CSF). IEEE, 263–275.
[4] Junxu Liu, Jian Lou, Li Xiong, Jinfei Liu, and Xiaofeng Meng. 2024. Cross-silo federated learning with record-level personalized differential privacy. In Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security. 303–317.
[5] Ken Liu, Shengyuan Hu, Steven Z Wu, and Virginia Smith. 2022. On privacy and personalization in cross-silo federated learning. Advances in neural information processing systems 35 (2022), 5925–5940.