Advancing Personalized Federated Learning: Group Privacy, Fairness, and Beyond

Filippo Galli, Gangsoo Zeong, Sayan Biswas, Catuscia Palamidessi, Tommaso Cucinotta

École Polytechnique & Inria

November 29, 2023

Overview

- Introduce two problems in federated learning
  - Non-IID and privacy
- Personalized federated learning with *d*-privacy
- Fairness aspect of personalized federated learning

Federated learning

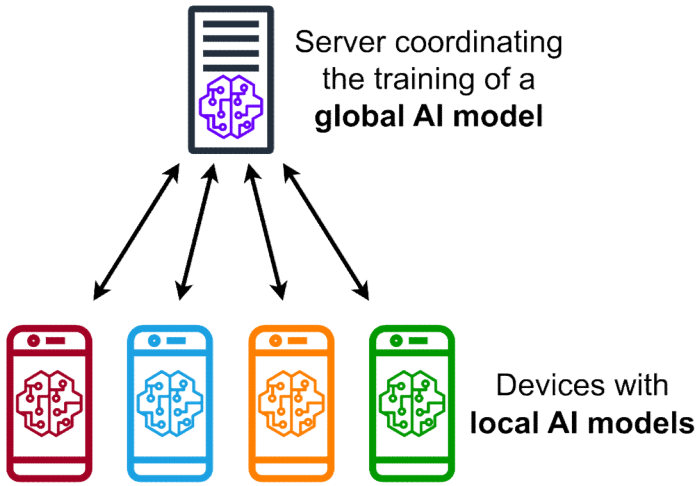

- Federated learning (FL) trains a machine learning model via multiple independent participants, each using its own dataset
- Training local models on local data samples and exchanging parameters to generate a global model shared by all nodes

But still, problem is not solved

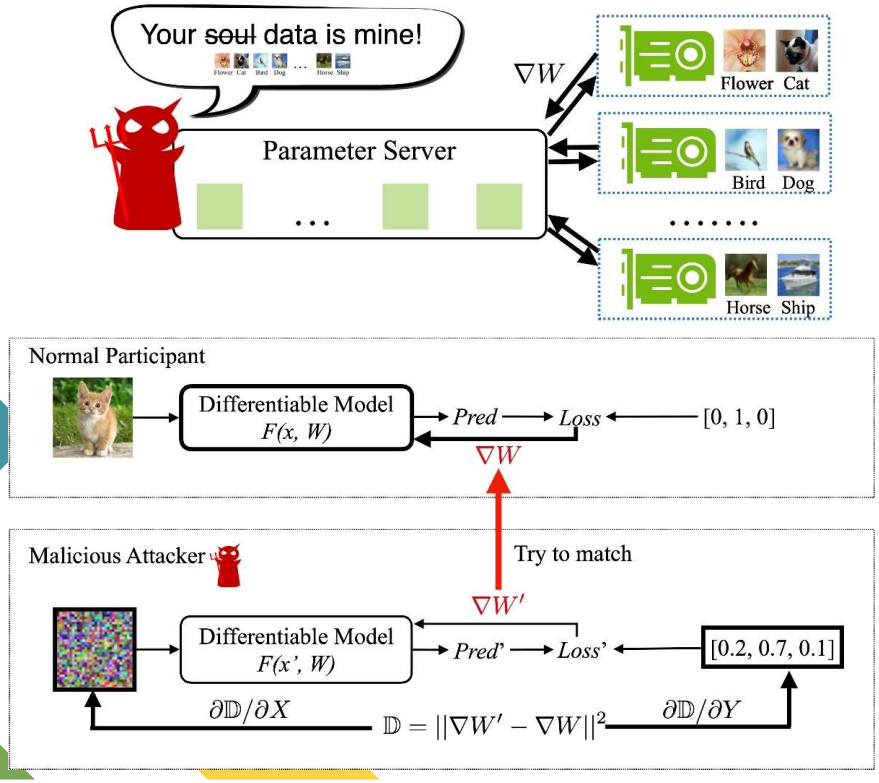

"Deep leakage from Gradients", NeurIPS 2019

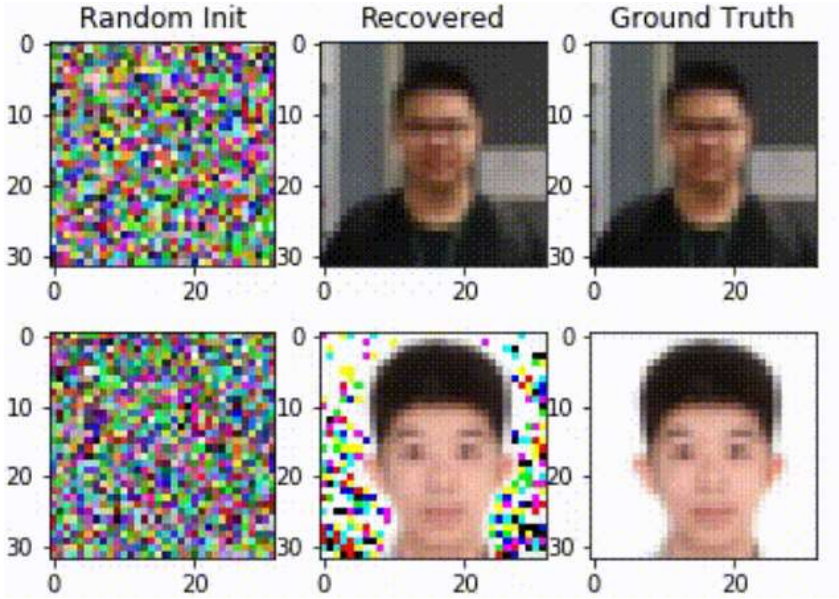

Differential privacy

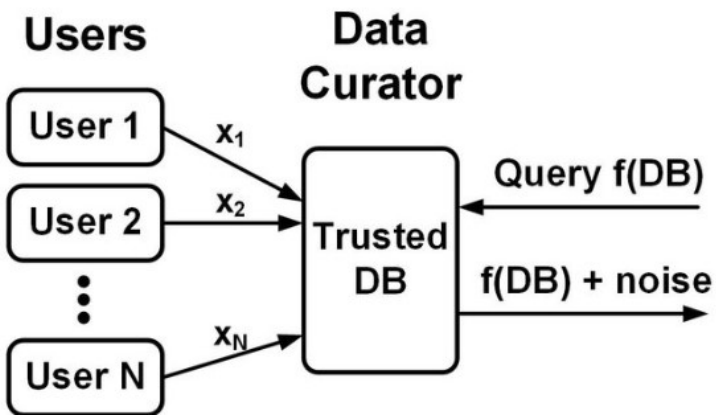

- "Deep Learning with Differential Privacy", CCS 2016
- Adding noise to model parameter(or Gradients)

Probability of seeing output O on input D_1

Probability of seeing output O on input D_2

$$\frac{\Pr[\mathcal{M}(D_1) \in O]}{\Pr[\mathcal{M}(D_2) \in O]} \leq e^{\varepsilon}$$

Indistinguishability: bounded ratio of probabilities

Personalized federated learning

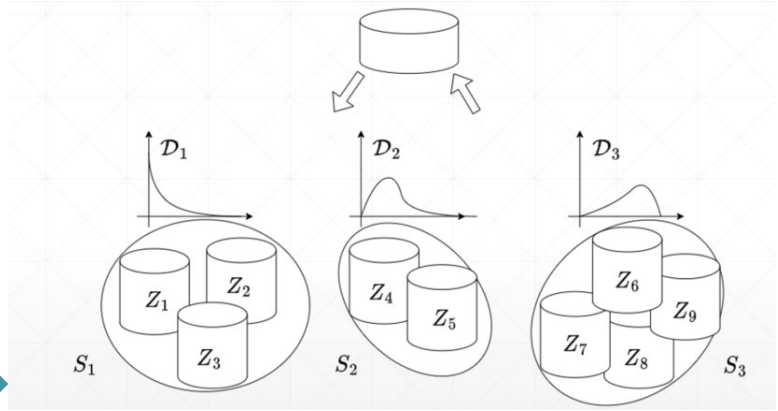

- Statistical challenges of federated learning
  - local data is non-IID
- Solution: Personalized federated learning

Problems

- Privacy
  - Avoiding the release of user's raw data is not enough
  - How can we enforce privacy?

- Personalization
  - How can we provide personalization?
  - Clients have different data distribution need different models. How do we adapt?

Our suggestion

- We propose an algorithm for personalized federated learning with local guarantees to provide privacy

- Our algorithm is motivated by the Iterative Federated Clustering Algorithm (IFCA) and builds on top of it to provide formal privacy guarantees

Overview of our algorithm

- IFCA with *d*-privacy

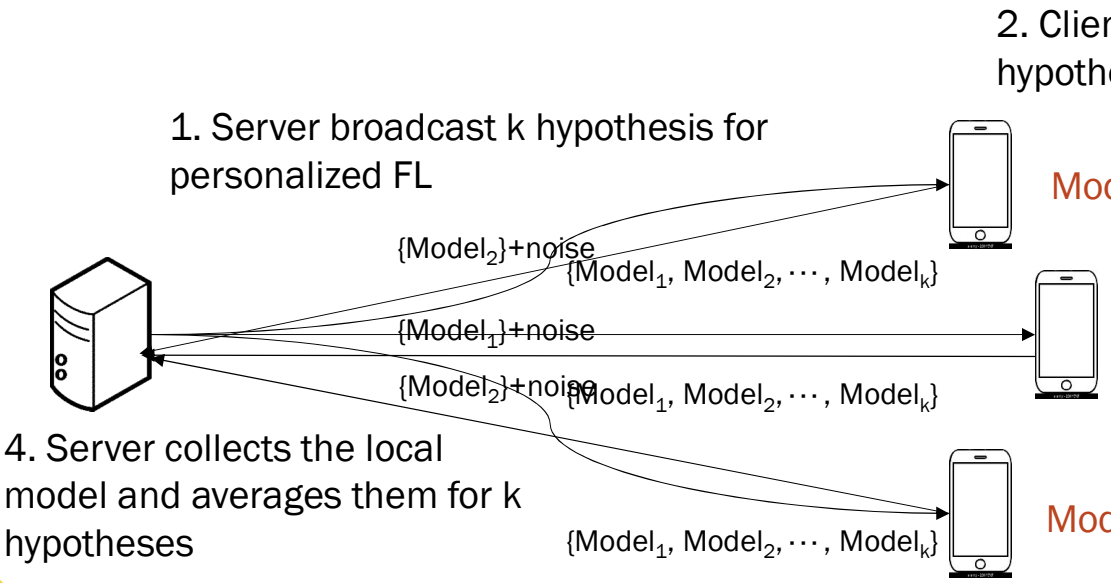

d-privacy(Metric privacy)

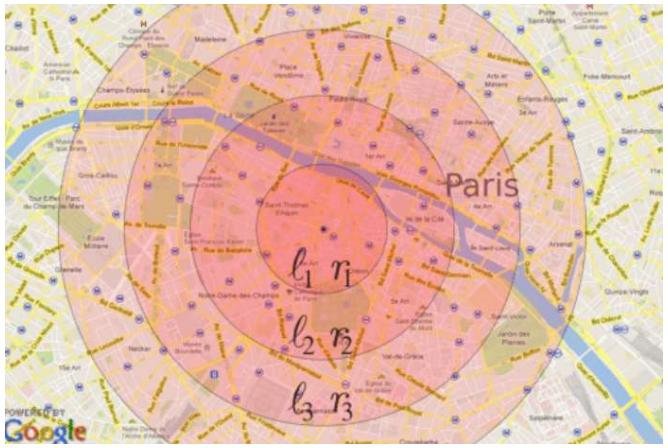

- Geo-Indistinguishability
  - A mechanism K satisfies ϵ-geo-indistinguishability iff for all x, x':

$$d_P(K(x), K(x')) \leq \epsilon d(x, x')$$

- There is an advantage of not having to do clipping

Algorithm for private and personalized federated learning

Algorithm 1 An algorithm for personalized federated learning with formal privacy guarantees in local neighborhoods.

Require: number of clusters k; initial hypotheses $\theta_j^{(0)}, j \in [k]$; number of rounds T; number of users per round U; number of local epochs E; local step size s; user batch size B_s; noise multiplier ν; local dataset Z_c held by user c.

1: **for** $t = \{0, 1, \ldots, T-1\}$ **do** ▷ Server-side loop
2: $C^{(t)} \leftarrow$ SampleUserSubset(U)
3: BroadcastParameterVectors($C^{(t)}; \theta_j^{(t)}, j \in [k]$)
4: **for** $c \in C^{(t)}$ **do** in parallel ▷ Client-side loop
5: $\bar{j} = \arg\min_{j \in [k]} F_c(\theta_j^{(t)}; Z_c)$
6: $\theta_{\bar{j},c}^{(t)} \leftarrow$ LocalUpdate($\theta_{\bar{j}}^{(t)}; s; E; Z_c$)
7: $\hat{\theta}_{\bar{j},c}^{(t)} \leftarrow$ SanitizeUpdate($\theta_{\bar{j},c}^{(t)}; \nu$)
8: **end for**
9: $\{S_1, \ldots, S_k\} =$ k-means($\hat{\theta}_{\bar{j},c}^{(t)}, c \in C^{(t)}; \theta_j^{(t)}, j \in [k]$)
10: $\theta_j^{(t+1)} \leftarrow \frac{1}{|S_j|} \sum_{c \in S_j} \hat{\theta}_{\bar{j},c}^{(t)}, \quad \forall j \in [k]$
11: **end for**

Algorithm 2 SanitizeUpdate obfuscates a vector $\theta \in \mathbb{R}^n$, with a Laplacian noise tuned on the radius of a certain neighborhood and centered around 0.

1: **function** SANITIZEUPDATE($\theta_{\bar{j}}^{(t)}; \theta_{\bar{j},c}^{(t)}; \nu$)
2: $\delta_c^{(t)} = \theta_{\bar{j},c}^{(t)} - \theta_{\bar{j}}^{(t)}$
3: $\varepsilon = \frac{n}{\nu \|\delta_c^{(t)}\|}$
4: Sample $\rho \sim \mathcal{L}_{0,\varepsilon}(x)$
5: $\hat{\theta}_{\bar{j},c}^{(t)} = \theta_{\bar{j},c}^{(t)} + \rho$
6: **return** $\hat{\theta}_{\bar{j},c}^{(t)}$
7: **end function**

Experimental result

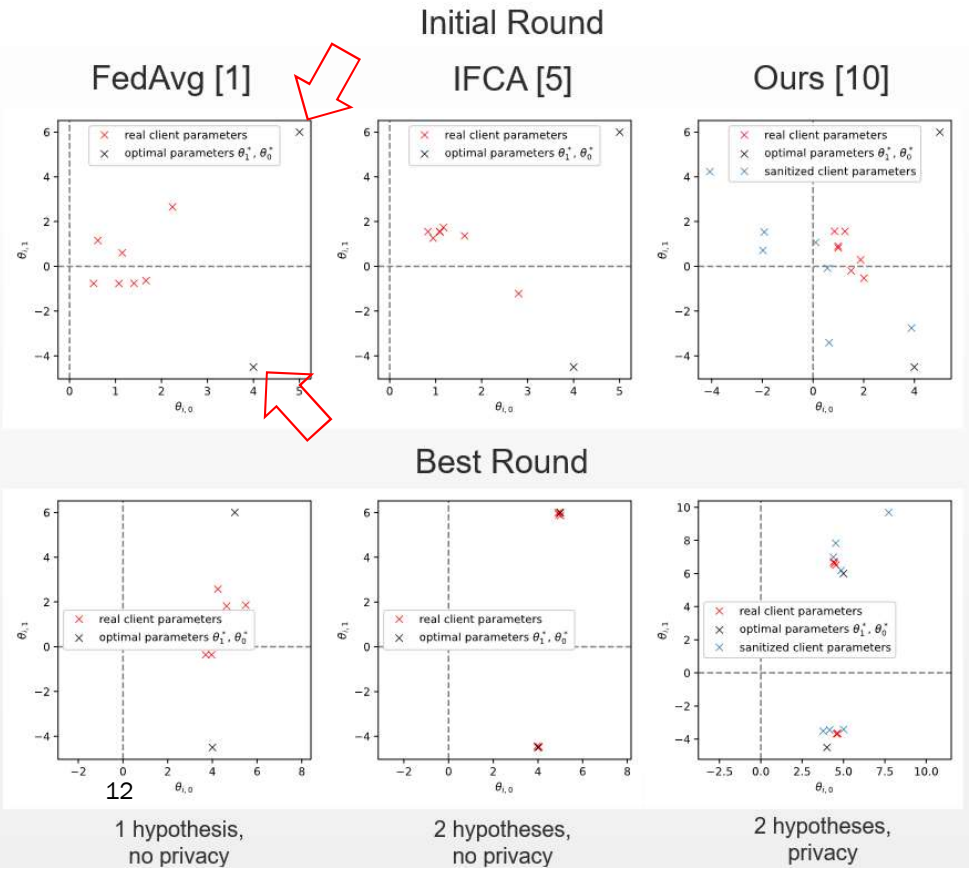

- Synthetic data
  - One hypothesis with non-sanitization
    - Failed to converge to the optimal parameter θ_1^* and θ_2^*
  - Two hypothesis with non-sanitization
    - Converge to the optimal parameter θ_1^* and θ_2^*
  - Two hypothesis with sanitization

Experimental result

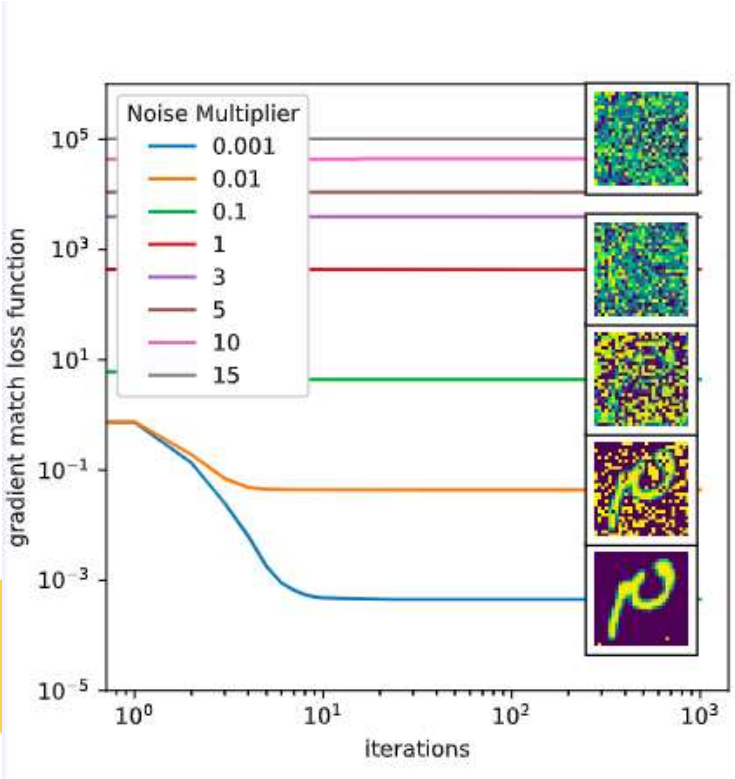

Noise multiplier	Average Accuracy
0	0.801
0.001	0.813
0.01	0.805
0.1	0.808
1	0.814
3	0.825
5	0.787
10	0.687
15	0.622

Fairness

- In addition to addressing privacy concerns, we also investigate the impact of our method on fairness in federated learning
- Several researches observed that systems aiming to protect privacy while ensuring fairness often involve a trade-off between the two
  - Privacy protection techniques based on differential privacy tend to minimize the impact of outliers or minorities within the overall dataset
- We presents experimental results demonstrating that the use of personalized FL improves fairness

Experimental result

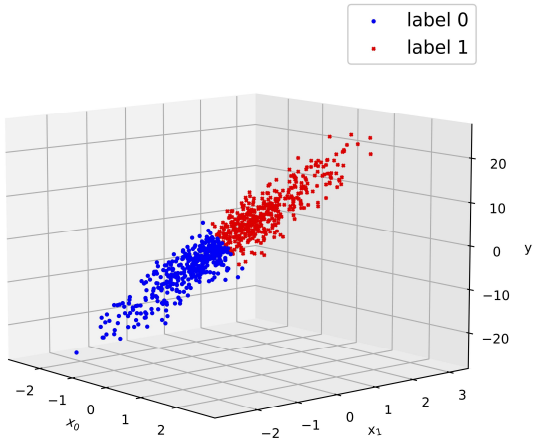

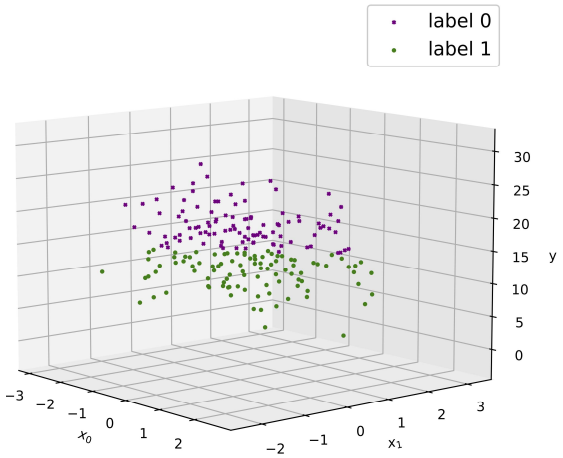

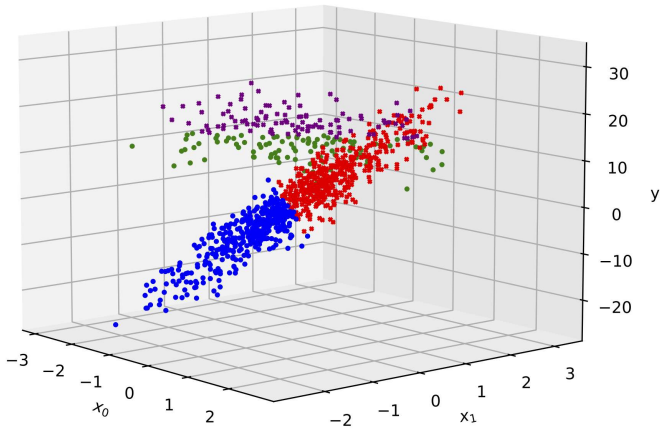

Experimental result

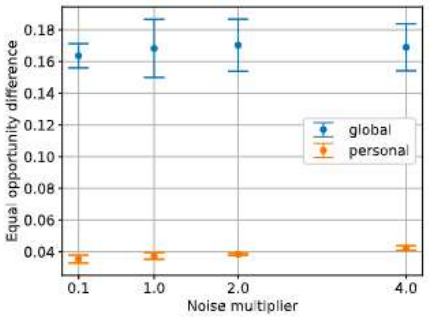

(a) Equal opportunity difference

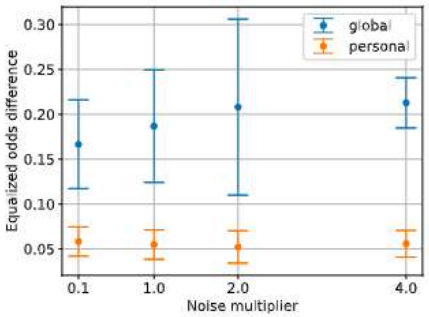

(b) Equalized odds difference

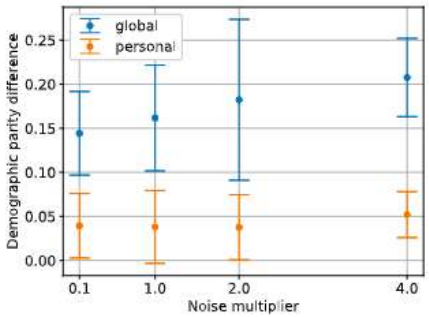

- Personalization of models enhances the group fairness under all the metrics
- A major reason behind this improvement of fairness is that the personalized model training optimizes for each group's data distribution

Conclusion

- We employ d-privacy techniques for sanitization on personalized federated learning
- Experimental results demonstrate the effectiveness of the proposed mechanism against the DLG attack
- We also evaluate the fairness of machine learning models trained using personalized federated learning and d-privacy

Q&A

Thanks you for listening!