Unified Privacy Guarantees for Decentralized Learning via Matrix Factorization

Dimitri Lerévérend[3]

Joint work with: Aurélien Bellet[1] Edwige Cyffers[2] Davide Frey[3] Romaric Gaudel[3] François Taïani[3]

February 3, 2026

[1]Inria, Université de Montpellier, INSERM, Montpellier, France
[2]Institute of Science and Technology Austria, Klosterneuburg, Austria
[3]Université de Rennes, Inria, CNRS, IRISA, Rennes, France

Decentralized Learning (DL)

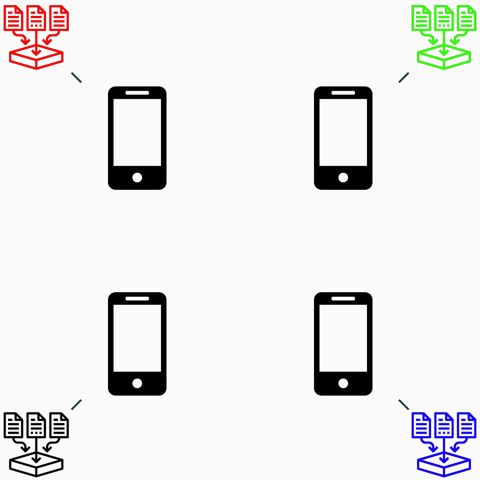

DL main caracteristics

- Possibly heterogeneous data

$\Longrightarrow$ How can we guarantee privacy in this setting?

Decentralized Learning (DL)

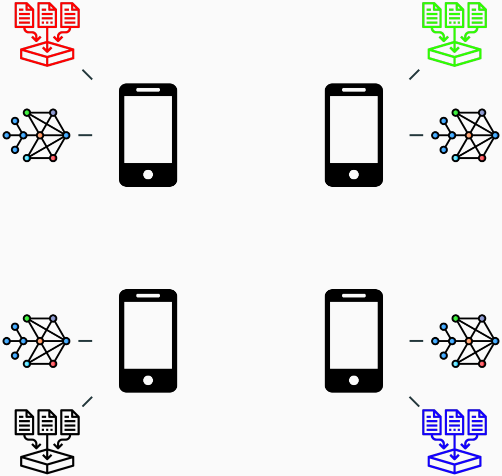

DL main caracteristics

- Possibly heterogeneous data
- Local model training

$\Longrightarrow$ How can we guarantee privacy in this setting?

Decentralized Learning (DL)

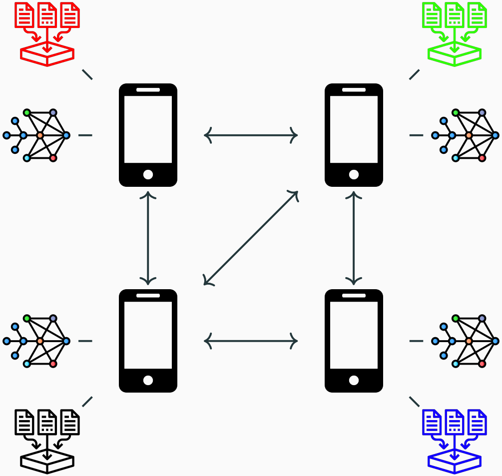

DL main caracteristics

- Possibly heterogeneous data
- Local model training
- Synchronous model exchanges
- Communication graph W
- No central server

$\Longrightarrow$ How can we guarantee privacy in this setting?

Differential Privacy in a nutshell

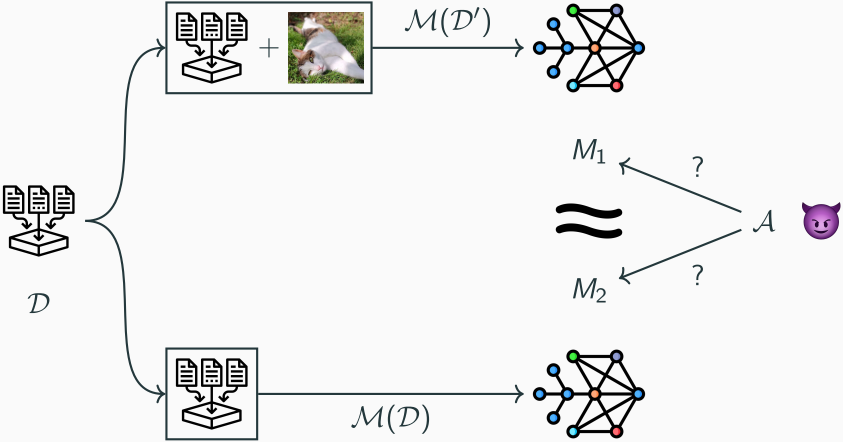

Differential Privacy in a nutshell

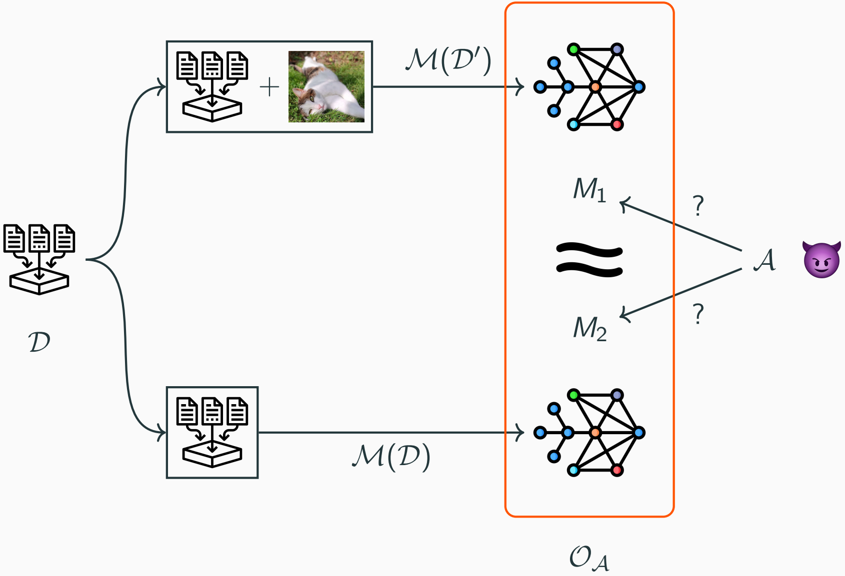

Who is the attacker?

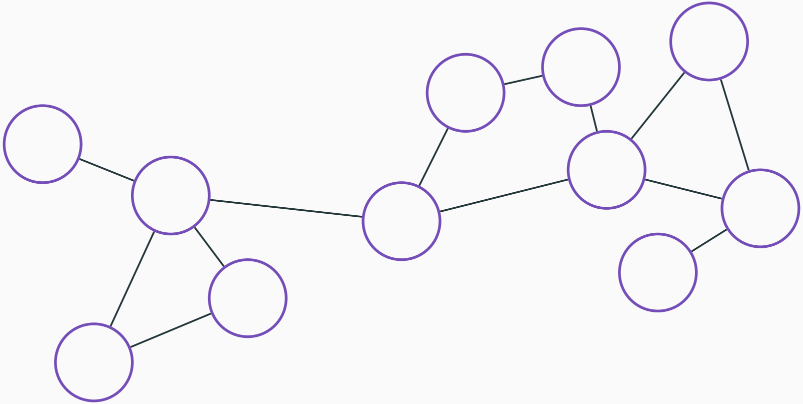

Observations of $\mathcal{A}$ ($\mathcal{O}_{\mathcal{A}}$)?

DP & DL: a game of observations

LDP

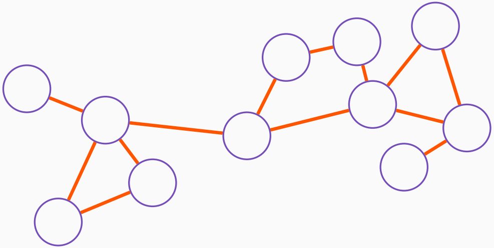

- $\mathcal{O}_{\mathcal{A}}$: all messages sent on the network

PNDP [3]

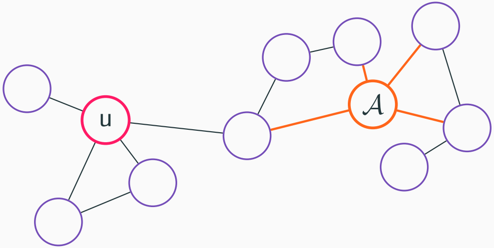

- $\mathcal{O}_{\mathcal{A}}$: messages received by node $\mathcal{A}$

Noisy approaches — DP-D-SGD

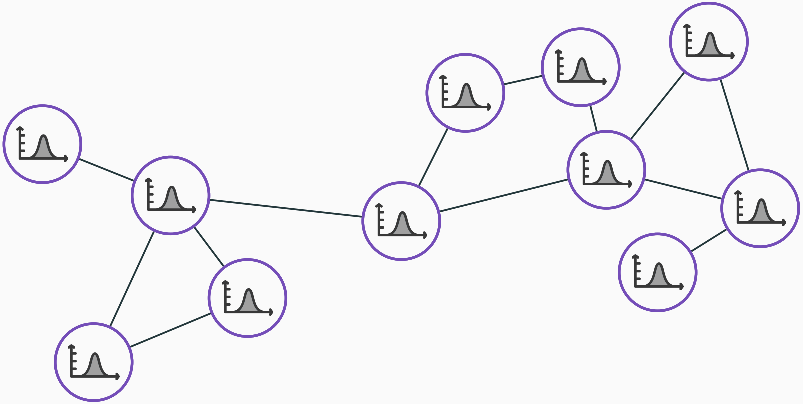

Limitations

- Impacts utility
- Difficult to scale

Correlated-noise approaches

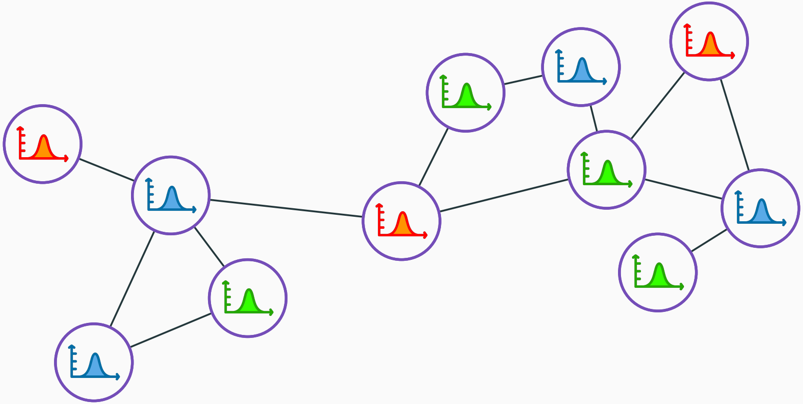

Correlation axis

- Space (Decor [1], ZIP-DL (ours, [2])) $\implies$ Complex analysis
- Time? $\implies$ No composition theorem.

Background: Matrix factorization in Centralized settings

Matrix Factorization in Centralized settings [6]

- Stack gradients and models into vectors:

$$A = \mathbf{1}_{i \geq j} = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 1 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & 1 \end{pmatrix}, \quad G = \begin{pmatrix} g_1 \\ g_2 \\ \vdots \\ g_t \end{pmatrix}, \quad \theta = \begin{pmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_t \end{pmatrix}$$

- A: workload matrix.
- We can rewrite SGD as a linear system:

$$\theta = 1_t \otimes \theta_0 - \eta A G.$$

Equivalent mechanism

$$\text{SGD: } \mathcal{M}(G) = AG$$

$$\text{DP-SGD: } \mathcal{M}(G) = A(G + Z), \quad Z \sim \mathcal{N}(0, \nu^2 I_t)$$

$$\text{Matrix mechanism: } \mathcal{M}(G) = AG + BZ, \quad Z \sim \mathcal{N}(0, \nu^2 I_t)$$

Matrix Factorization in Centralized settings [6]

Goal: find good factorizations $A = BC$.

Theorem — DP guarantees for Centralized learning [4]

- Hypothesis:
  - Centralized/Federated setting
  - $A = BC$ and A is squared & lower triangular & invertible.
- Then, $\mathcal{M}(G) = B(CG + Z)$ with $Z \sim \mathcal{N}(0, \nu^2)$ with $\nu = \sigma\,\mathrm{sens}(C)$ is $\frac{1}{\sigma}$-GDP, even under adaptive G.

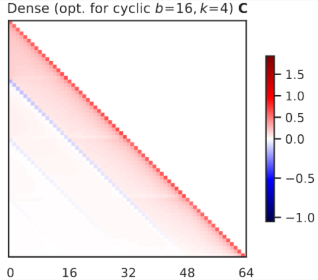

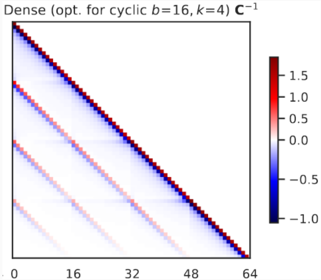

Matrix Factorization in Centralized settings [6]

Goal: find good factorizations $A = BC$.

Theorem — DP guarantees for Centralized learning [4]

- Hypothesis:
- -GDP,

Our objectives

1. Adapt the matrix-factorization formalism to decentralized settings.
2. Extend the centralized theorem by relaxing structural assumptions.
3. Derive tighter privacy accounting for decentralized mechanisms.
4. Introduce MAFALDA-SGD for optimized correlated noise.

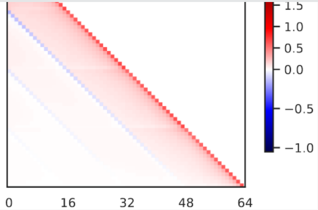

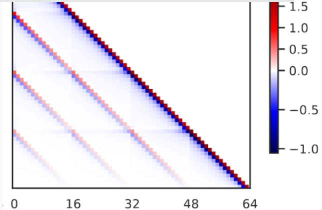

Our work: Unifying it all in DL

Adapting MF to DL

We stack through both time and space: T block of n values, one for each node.

Communication workload

$$\mathbf{W}_T = \begin{bmatrix} I_n & 0 & 0 & \dots & 0 \\ W & I_n & 0 & \dots & 0 \\ W^2 & W & I_n & \dots & 0 \\ \dots & \dots & \dots & \ddots & \dots \\ W^{T-1} & W^{T-2} & W^{T-3} & \dots & I_n \end{bmatrix}, \quad G = \begin{bmatrix} G_1 \\ G_2 \\ G_3 \\ \vdots \\ G_T \end{bmatrix}$$

- W: communication matrix

Matrix Factorization in DL — high level view

$$
\begin{array}{r}\text{Iter } 1\{\\ \text{Iter } 2\{\\ \\ \text{Iter } T\{\end{array}
\underbrace{\begin{pmatrix}\theta_1\\ \theta_2\\ \vdots\\ \theta_T\end{pmatrix}}_{\text{Stacked models}}
= \overbrace{(I_T \otimes W)}^{\text{Local averaging with neighbors}}
\overbrace{\left(\underbrace{\overbrace{\begin{pmatrix} I_n\\ W\\ \vdots\\ W^{T-1}\end{pmatrix}}^{M}}_{\text{Data Independant}}\theta_0 - \eta \underbrace{\overbrace{\begin{bmatrix} I_n & & & \\ W & \ddots & & \\ \vdots & \ddots & \ddots & \\ W^{T-1} & \dots & W & I_n\end{bmatrix}}^{\text{Workload } W_T}\left(\begin{pmatrix}G_0\\ G_1\\ \vdots\\ G_{T-1}\end{pmatrix} + C^{\dagger} Z\right)}_{\mathcal{O}_{\mathcal{A}}}\right)}^{\text{Messages}}
$$

Gossip matrix $\in \mathbb{R}^{n \times n}$

Initial local models $\in \mathbb{R}^{n \times d}$

Drawn from $\mathcal{N}(0, \sigma^2 \text{sens}(c)^2 I_{nT})$

Optimized under constraints

Attacker observations:

$$\mathcal{O}_{\mathcal{A}} = AG + BZ$$

- A is a rectangular matrix.
- A has a column-echelon structure.

Adaptive privacy guarantees

Theorem — Unified DP guarantees for DL

- Hypothesis:
  - Decentralized learning settings
  - $A = BC$ and A is rectangular & column echelon.
- Then, $\mathcal{M}(G) = B(CG + Z)$ with $Z \sim \mathcal{N}\left(0, \nu^2\right)$ with $\nu = \sigma\,\mathrm{sens}(C; B)$ is $\frac{1}{\sigma}$-GDP, even under adaptive G.
- $\mathrm{sens}_\Pi(C; B) \leq \max_{\pi \in \Pi} \sum_{s,t \in \pi} \left| \left(C^\top B^\dagger B C\right)_{s,t} \right|$

Novelties

- Wider range of workloads
- Extends the notion of sensitivity

Remark — Recovering known threat models

LDP + DP-D-SGD

- Attacker observes all noisy gradients.
- IID noises accross all nodes and rounds.
- $A = B = \mathbf{W}_T,\ C = I_{nT}$

PNDP [3]

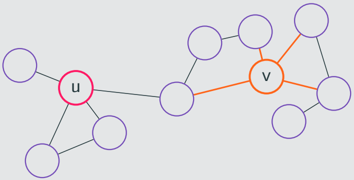

- Attacker $\mathcal{A}$ observes a subset of noisy gradients
- $P_{\mathcal{A}}\mathbf{W}_T G$ projection on the gradients observed by $\mathcal{A}$

Application 1 — Tighter accounting

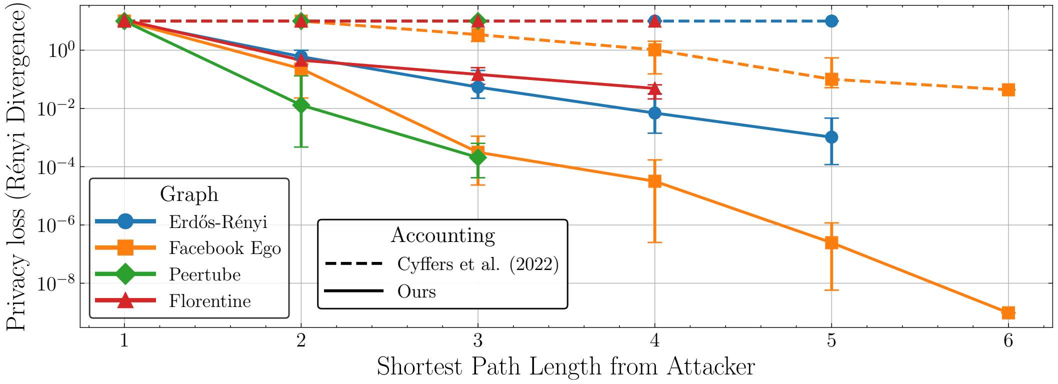

- Recover existing privacy accounting such as PNDP [3].
- We derive tighter PNDP bounds for DP-D-SGD.

Application 2: Correlation optimization (Mafalda-SGD)

MAtrix FActorization for Local Differential PrivAte-SGD (MAFALDA-SGD)

- Adapt optimization objective to LDP setting.
- Force same noise pattern for all nodes.
- The new minimization problem is

$$\mathcal{L}_{\text{opti}}\left(\mathbf{W}_T, C_{\text{local}}\right) = \underset{\Pi_{\text{local}}}{\text{sens}}\left(C_{\text{local}}\right)^2 \left\| L C_{\text{local}}^{\dagger} \right\|_F^2$$

with L the Choleski decomposition such that

$$L^\top L = \sum_{i=1}^{n} A_i^\top A_i,$$

$$A_i := \left[(I_T \otimes W)\,\mathbf{W}_T \mathbf{K}^{(T,n)}\right]_{[:,iT:(i+1)T-1]}$$

- Solve $\min_{C_{\text{local}}} \mathcal{L}_{\text{opti}}\left(\mathbf{W}_T, C_{\text{local}}\right)$ using L-BFGS [6].

Experimental results — FEMNIST

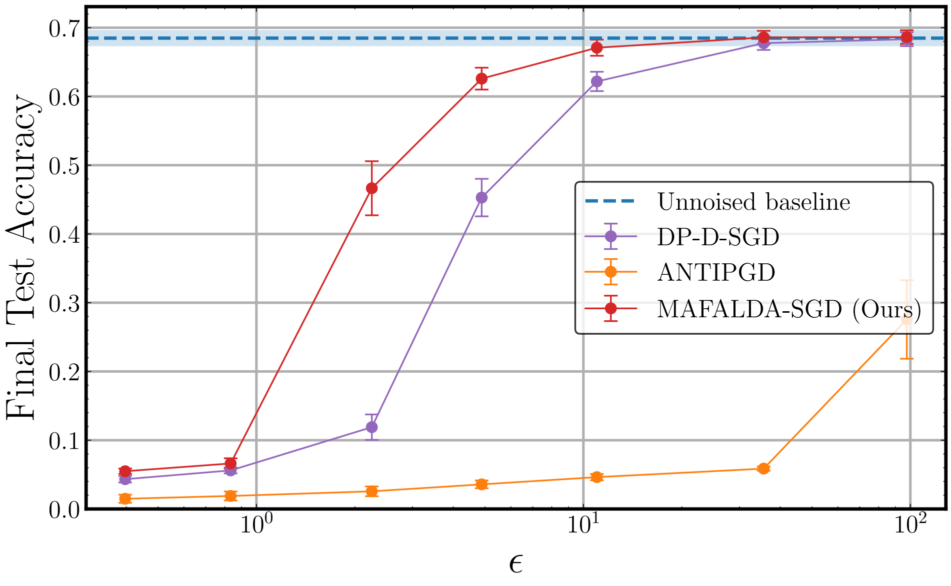

Conclusion

Our work

- Unifies DP guarantees under various noise patterns/attackers
- Derives tighter privacy guarantees for MF mechanisms
- Introduces a novel algorithm that outperforms LDP baselines

Future works

- Explore localized optimums and other threat models
- Find cross-nodes optimal correlations

References i

- Youssef Allouah, Anastasia Koloskova, Aymane El Firdoussi, Martin Jaggi, and Rachid Guerraoui.
  The privacy power of correlated noise in decentralized learning.
  In *Proceedings of the 41st International Conference on Machine Learning*, pages 1115–1143, 2024.
- Sayan Biswas, Davide Frey, Romaric Gaudel, Anne-Marie Kermarrec, Dimitri Lerévérend, Rafael Pires, Rishi Sharma, and François Taïani.
  Low-cost privacy-preserving decentralized learning.
  Proceedings on Privacy Enhancing Technologies, 2025.

Edwige Cyffers, Mathieu Even, Aurélien Bellet, and Laurent Massoulié.
Muffliato: Peer-to-Peer Privacy Amplification for Decentralized Optimization and Averaging.
NeurIPS, 2022.

Sergey Denisov, H Brendan McMahan, John Rush, Adam Smith, and Abhradeep Guha Thakurta.
Improved differential privacy for sgd via optimal private linear operators on adaptive streams.
Advances in Neural Information Processing Systems, 35:5910–5924, 2022.

Abdellah El Mrini, Edwige Cyffers, and Aurélien Bellet.
Privacy attacks in decentralized learning.
ICML'24. JMLR.org, 2024.

Krishna Pillutla, Jalaj Upadhyay, Christopher A. Choquette-Choo, Krishnamurthy Dvijotham, Arun Ganesh, Monika Henzinger, Jonathan Katz, Ryan McKenna, H. Brendan McMahan, Keith Rush, Thomas Steinke, and Abhradeep Thakurta.
Correlated noise mechanisms for differentially private learning, 2025.

Experimental results — Housing