

# Differentially Private Federated Learning with Time-Adaptive Privacy Spending

International Conference on Learning Representations, Singapore

Poster Session 5, Apr 25, 2025

Shahrzad  
Kianidehkordi



Nupur  
Kulkarni



Adam  
Dziedzic



Stark C.  
Draper



Franziska  
Boenisch



The Edward S. Rogers Sr. Department of  
Electrical & Computer Engineering  
**UNIVERSITY OF TORONTO**



**CISPA**  
HELMHOLTZ-ZENTRUM FÜR  
INFORMATIONSSICHERHEIT



**Mitacs**  
**Globalink**  
Research Award  
Abroad




**NSERC CGS D3**  
**& Discovery**  
**Research Grant**




**DiDi**  
Graduate  
Award

Distribute training tasks across a group of “honest-but-curious” clients



**Centralized learning**

vs.




**Federated learning**

How to keep data private?

**Keep data decentralized by  
federated learning [1]**

But clients privacy can leak via sharing local model updates!



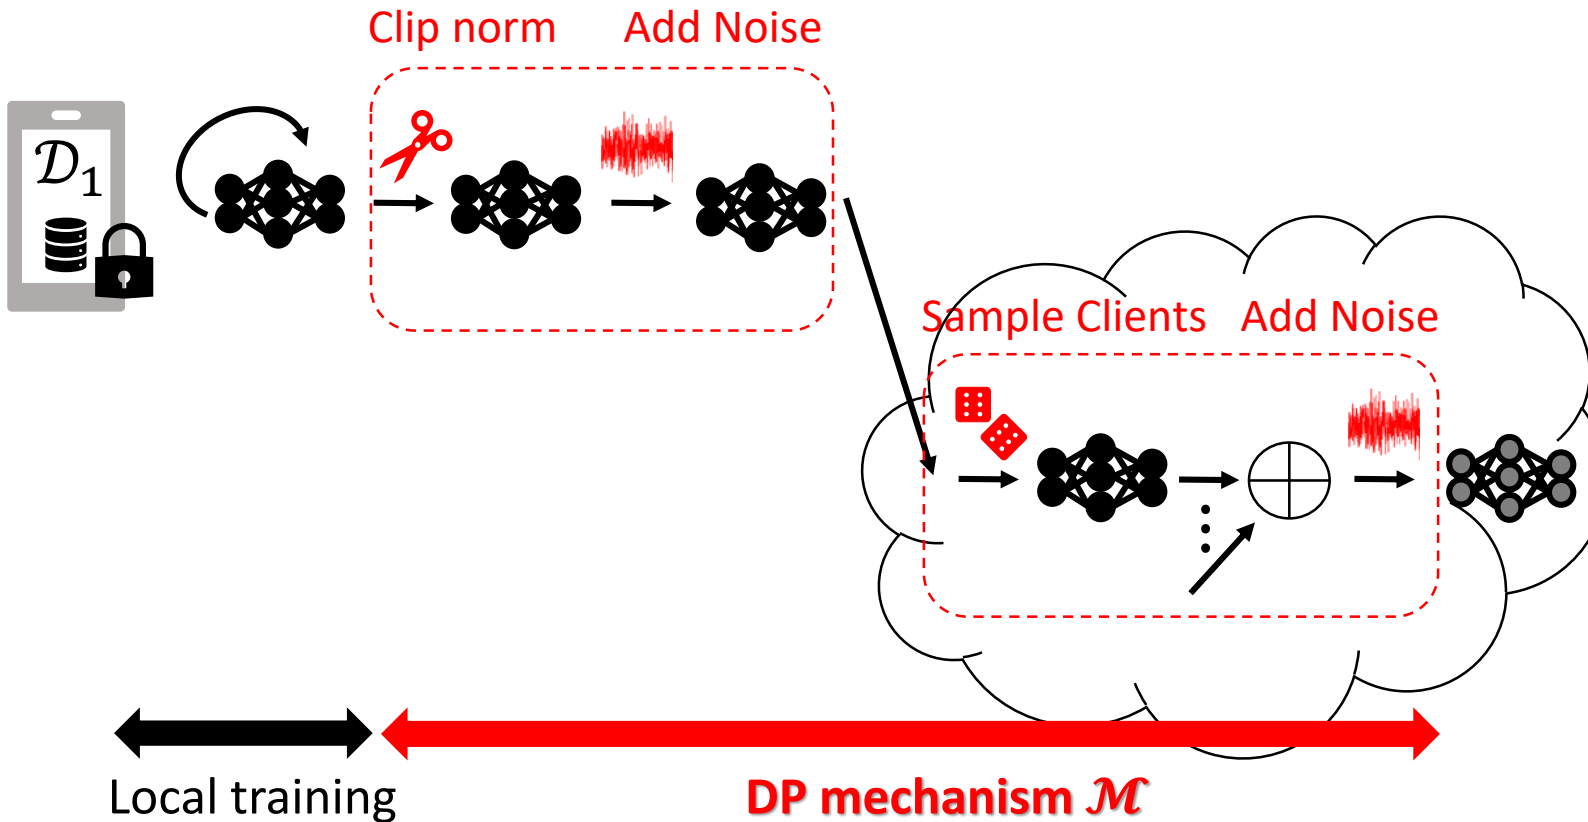
How to keep data private?

Keep data decentralized by federated learning (FL) [1]

**But FL is not enough for privacy!**


e.g., “**membership inference attack**” [2].

# Add controlled perturbation via differential privacy mechanism




Output distributions of a **DP mechanism  $\mathcal{M}$**  on any neighboring data  $\mathcal{D} \approx \mathcal{D}'$  is nearly **indistinguishable!**

$$\frac{\mathcal{P}(\mathcal{M}(\mathcal{D}) = \theta)}{\mathcal{P}(\mathcal{M}(\mathcal{D}') = \theta)} \leq e^\epsilon$$




# Fixing privacy budget expenditure over time is not necessarily optimal!



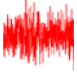





- Fixing privacy expenditure → Often fixes DP perturbation (noise)
- Convergence → Over time, gradient norms (signal) decreases

➔ **Poor SNR in later rounds!**

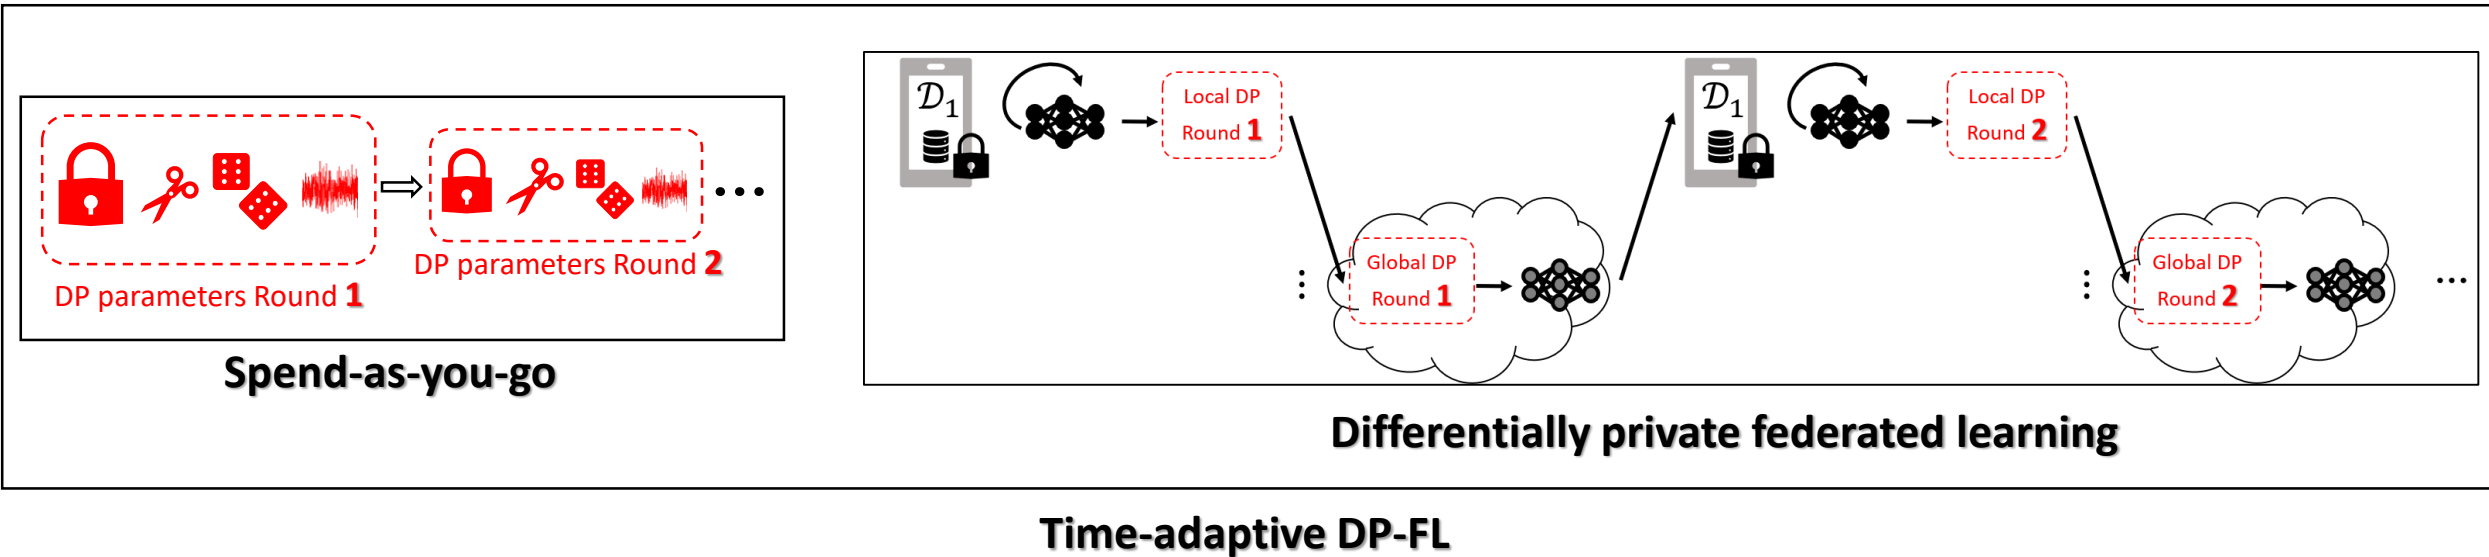
# Our time-adaptive DP-FL spends privacy non-uniformly over training!



- In early rounds → save privacy budget via stricter  → sample more strictly  → noise increases 
- In later rounds → stop saving → more relaxed  → sample less strictly  → noise reduces 

➡ **SNR improves in later iterations!**

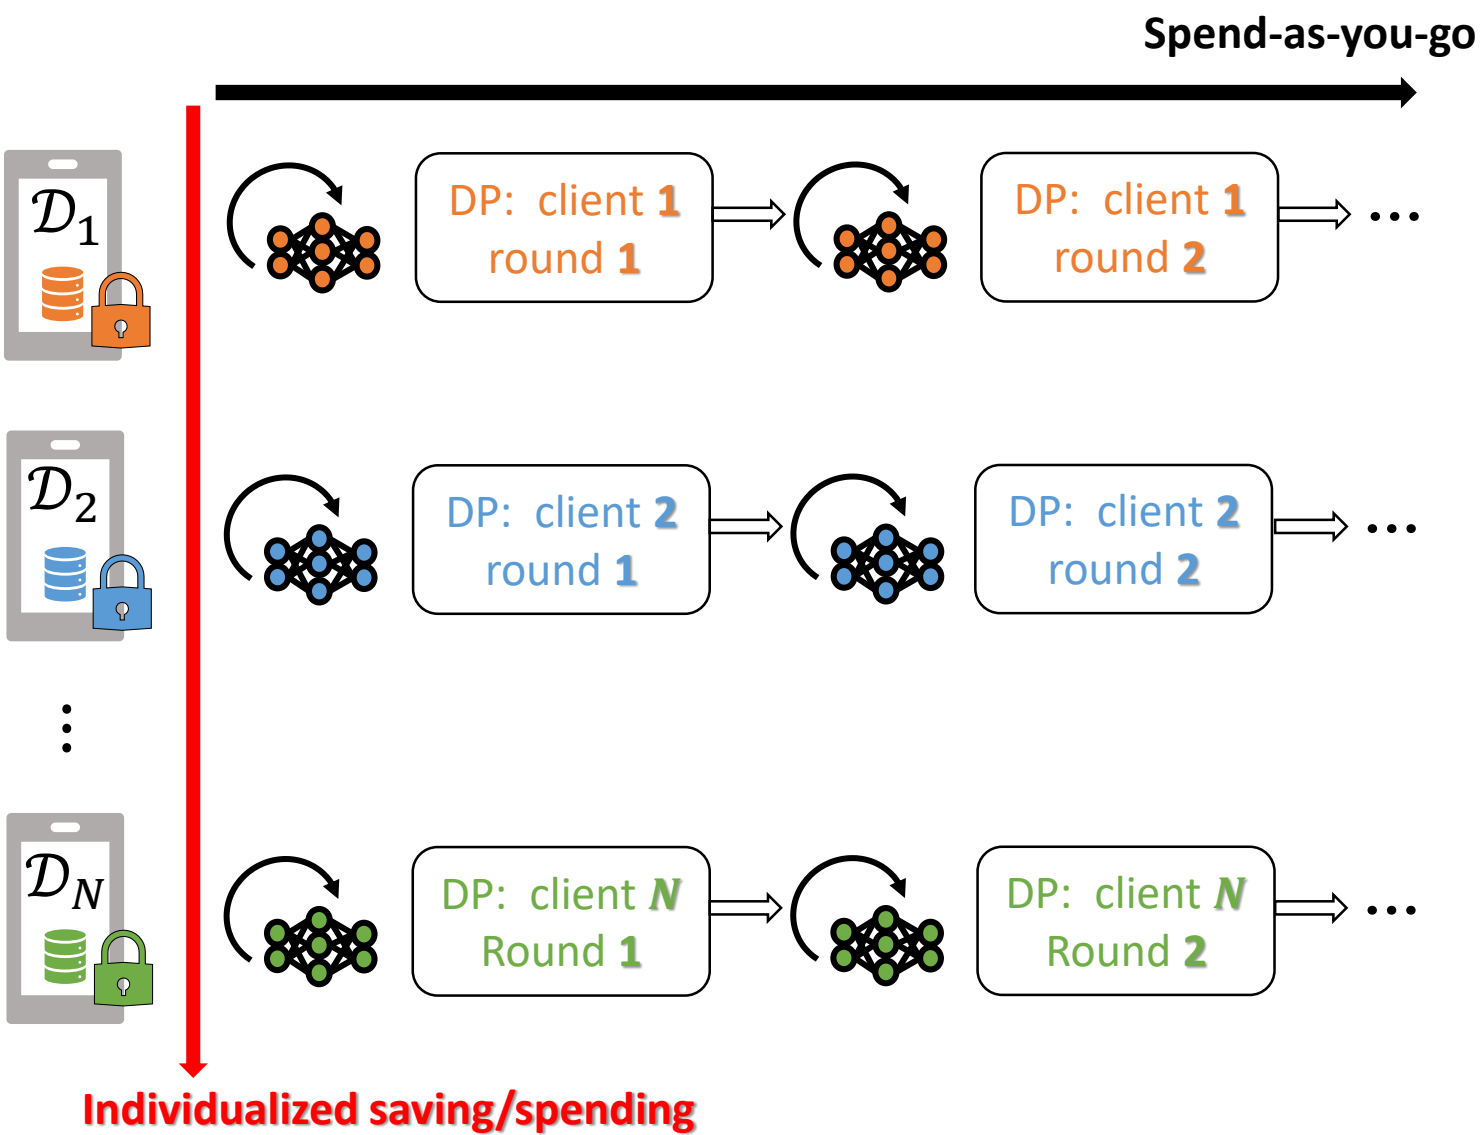
We separate spend-as-you-go module and DP-FL module!



➤ If set privacy parameters dependent on gradient norms (signal) → Additional privacy risk!


➡ **We set privacy parameters dependent on only clients' privacy constraints, and not training signals!**

# Our time-adaptive DP-FL spends privacy non-uniformly amongst clients!



➤ We let clients with smaller privacy budgets spend more non-uniformly!


➡ **Reduces clipping bias in expectation!**






# Saving budgets in initial rounds & spending more later improve utility!

- In the first half of training, we save privacy budgets to spend more in the second half of training.
- Time adaptive DP-FL improves accuracy in later rounds compared to Individualized DP-FedAvg. 😊
- Both adhere to privacy budgets. 😊  
TA-DP-FL follows spend-as-you-go.  
IDP-FedAvg spends budget uniformly.



# Thank you



TA-DP-FL lets clients save privacy budgets for rounds requiring higher SNR.

Spend-as-you-go

DP-FL

TA-DP-FL

TA-DP-FL's privacy module is data-independent, relies only on individual privacy budgets, enables stricter-budget clients to spend less uniformly.

## Questions?

Paper: <https://arxiv.org/pdf/2502.18706>

Emails:

shahrzad.kianidehkordi@mail.utoronto.ca,  
nuku00001@stud.uni-saarland.de,  
adam.dziedzic@cispa.de,  
stark.draper@utoronto.ca,  
boenisch@cispa.de

**Open problems:** Adapt privacy spending parameters to data while preserving privacy. Combine privacy spending with adaptive clipping techniques.