Georgia Tech School of Computer Science
College of Computing

LDP-FED: FEDERATED LEARNING WITH LOCAL DIFFERENTIAL PRIVACY

STACEY TRUEX, LING LIU, KA-HO CHOW, MEHMET EMRE GURSOY, AND WENQI WEI

CREATING THE NEXT®

This work was performed during 2019 Fall and 2020 Spring semester. Stacey would like to acknowledge the 2019 - 2020 IBM PhD fellowship recently awarded to her and the support from IBM Almaden under the Enterprise AI, Systems & Solutions Research Group (director: Sandeep Gopisetty).

TALK OUTLINE

- Motivation
- Deep Learning & Federated Training
- Privacy Leakage in Federated Learning Systems
- Differential Privacy for Iterative ML Training Mechanisms
    - α-Condensed Local Differential Privacy
- LDP-Fed: Federated Learning with Local Differential Privacy
    - LDP Module, k-Selection Module
- LDP-Fed Performance & Features
- Conclusion

GROWTH OF DATA COLLECTION

In 2020 there will be 40x more bytes of data than there are stars in the observable universe.

DOMO report

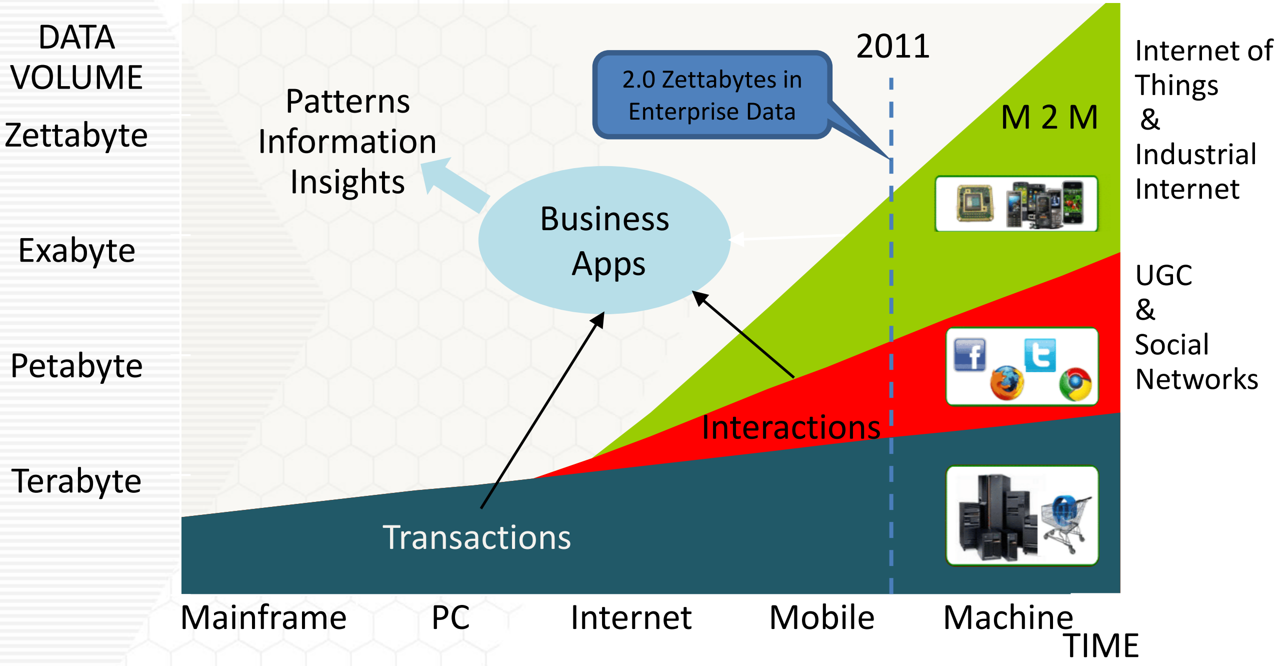

Infographic source: rightedge

GROWTH OF MACHINE LEARNING SERVICES

Machine learning

5 Year Growth Rate: 34%

- Published patent applications for Patent Classification G06N “Computer Systems Based on Specific Computational Models” grew at a compound annual rate of 34% from 2013 to 2017.
- This includes machine learning and artificial neural networks.

Forbes article

Company	2017 Published Applications
IBM	654
Microsoft	139
Google	127
LinkedIn	70
Facebook	66
Intel	52
Fujitsu	49

Data Science platforms that support machine learning are predicted to grow at a 13% CAGR through 2021

REACTION: DEMAND FOR PRIVACY

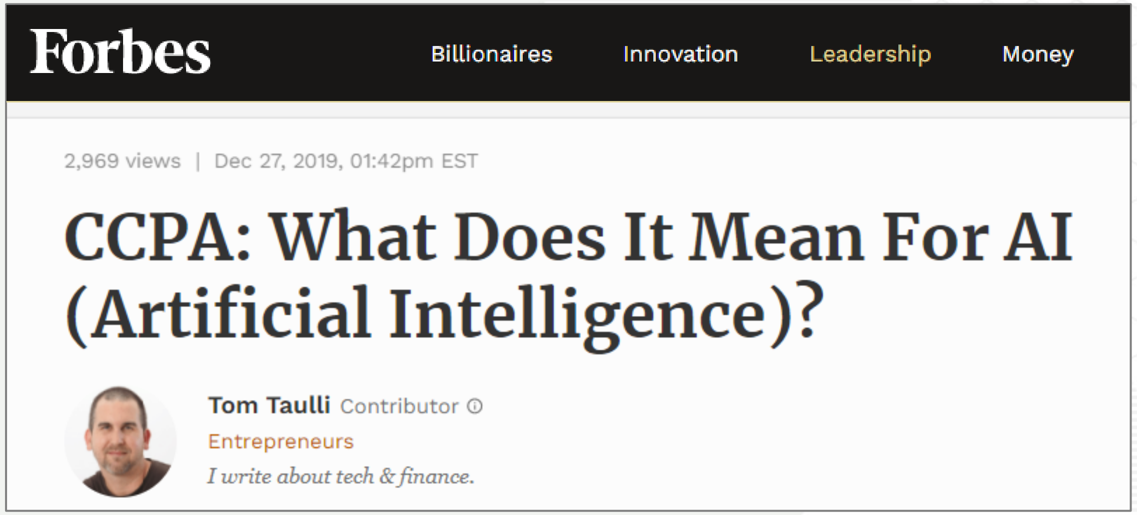

2,969 views | Dec 27, 2019, 01:42pm EST

CCPA: What Does It Mean For AI (Artificial Intelligence)?

Tom Taulli Contributor
Entrepreneurs
I write about tech & finance.

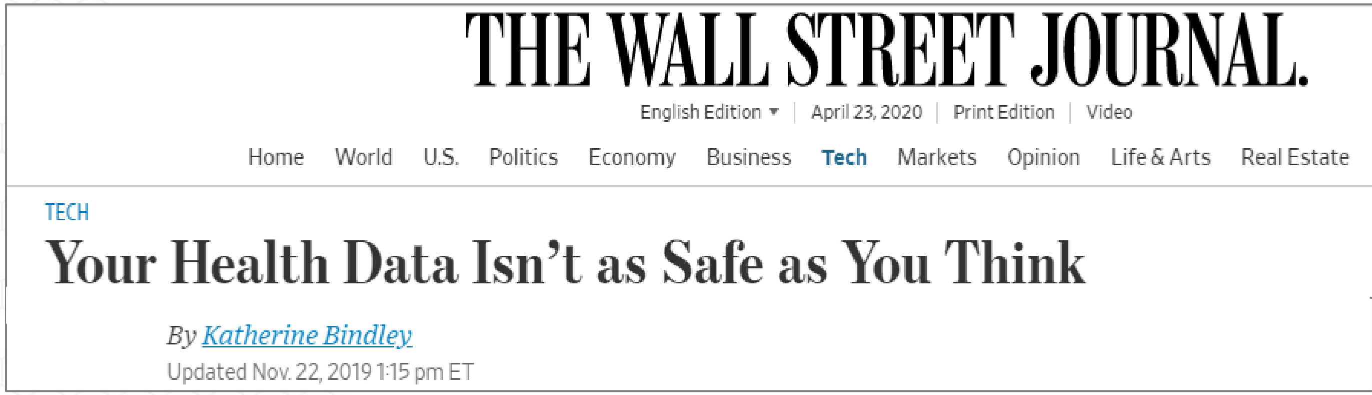

TECH

Your Health Data Isn't as Safe as You Think

By Katherine Bindley
Updated Nov. 22, 2019 1:15 pm ET

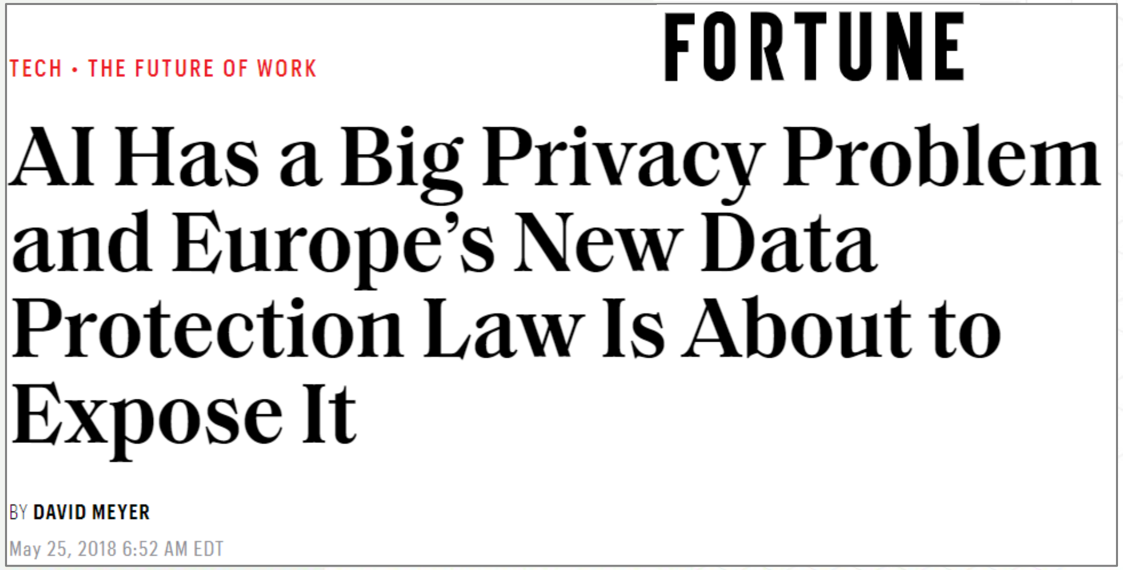

TECH • THE FUTURE OF WORK

AI Has a Big Privacy Problem and Europe's New Data Protection Law Is About to Expose It

BY **DAVID MEYER**
May 25, 2018 6:52 AM EDT

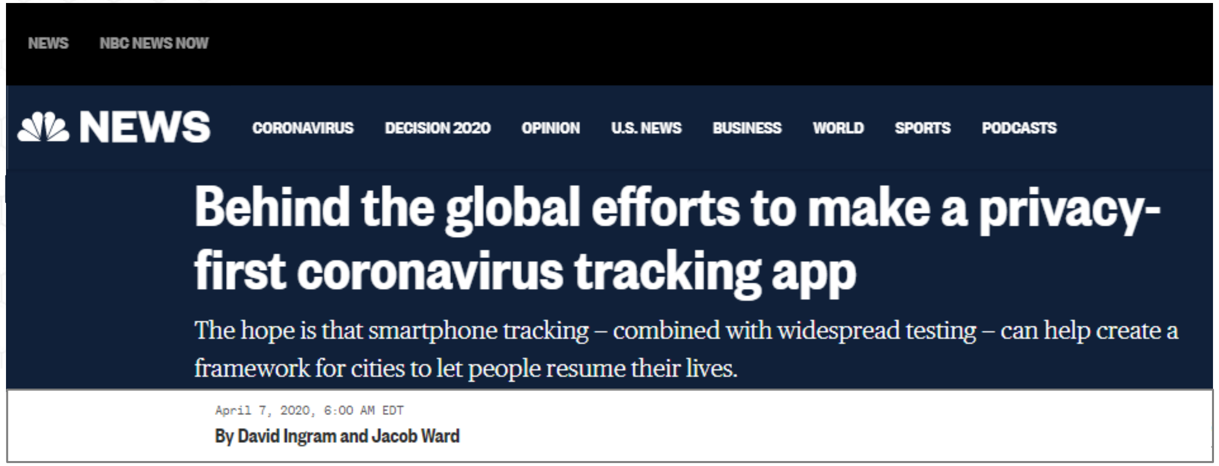

Behind the global efforts to make a privacy-first coronavirus tracking app

The hope is that smartphone tracking – combined with widespread testing – can help create a framework for cities to let people resume their lives.

April 7, 2020, 6:00 AM EDT
By David Ingram and Jacob Ward

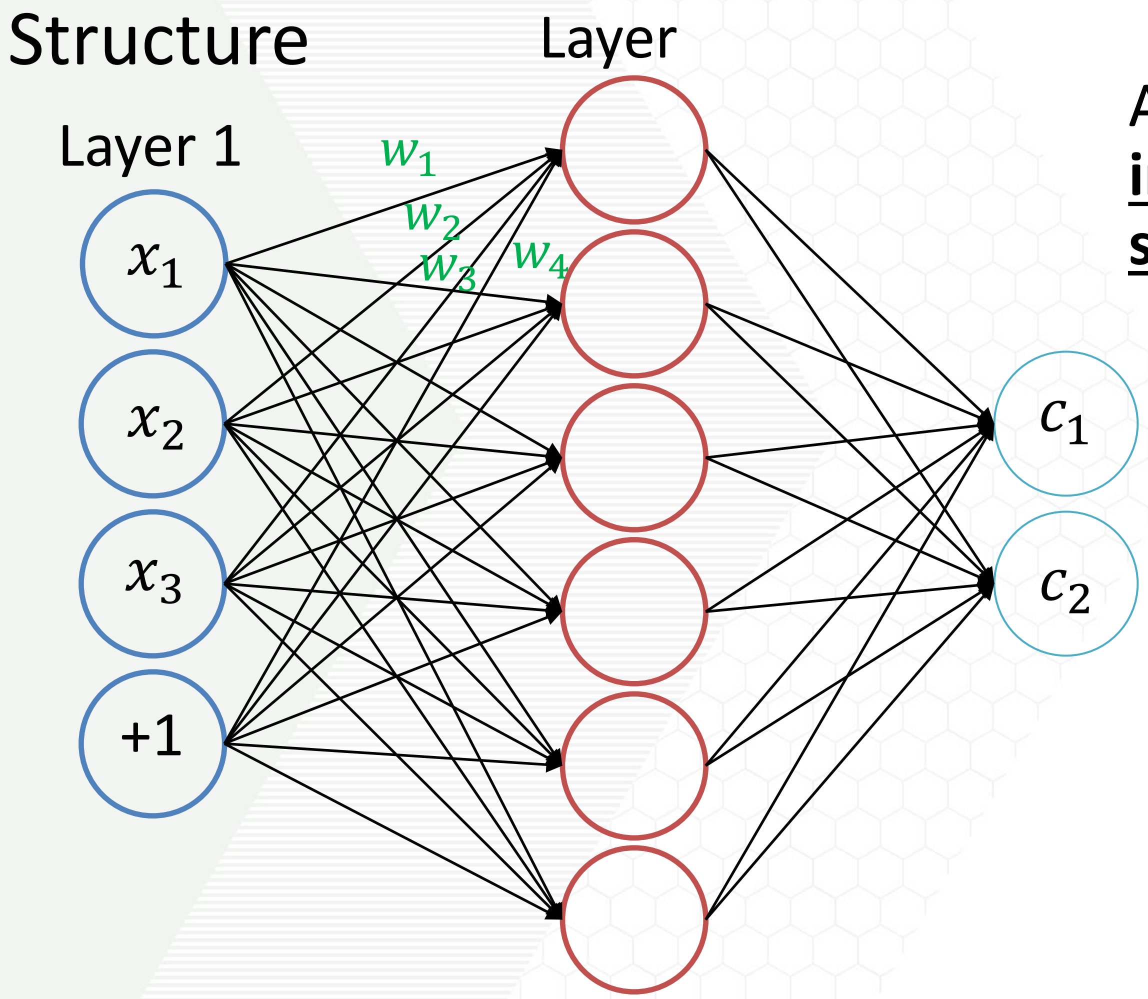

Activation function converts **input signals** to an **output signal**

An activation function is applied to the sum of the product of input signals and their **corresponding weights**

Learning Process

1. Shuffle data and divide into batches
2. Feed batches forward through the network
3. Calculate Error
4. Backpropagate the error
5. Use gradients to update weights

Multiple *epochs*

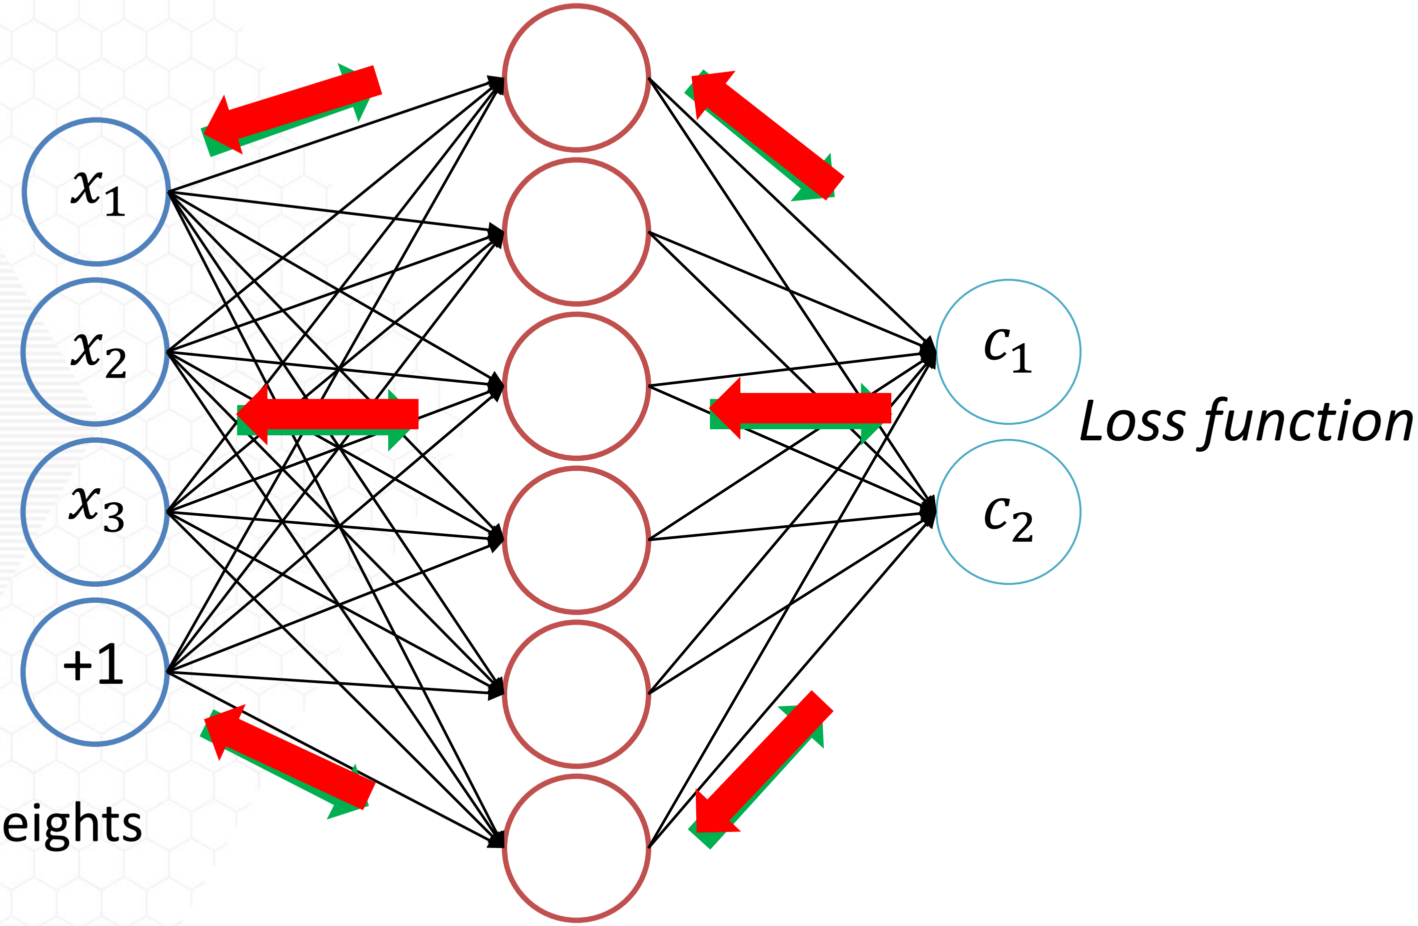

FEDERATED TRAINING: LEVERAGING EDGE DEVICES

1. Same model structure & parameters are initialized at each participant
2. Each participant conducts local training on their private dataset, resulting in updated parameters
3. Locally updated model parameters are sent to the parameter server
4. Server aggregates the parameter updates using Federated Averaging
5. New, aggregated parameters are broadcast to all participants

Multiple *iterations*

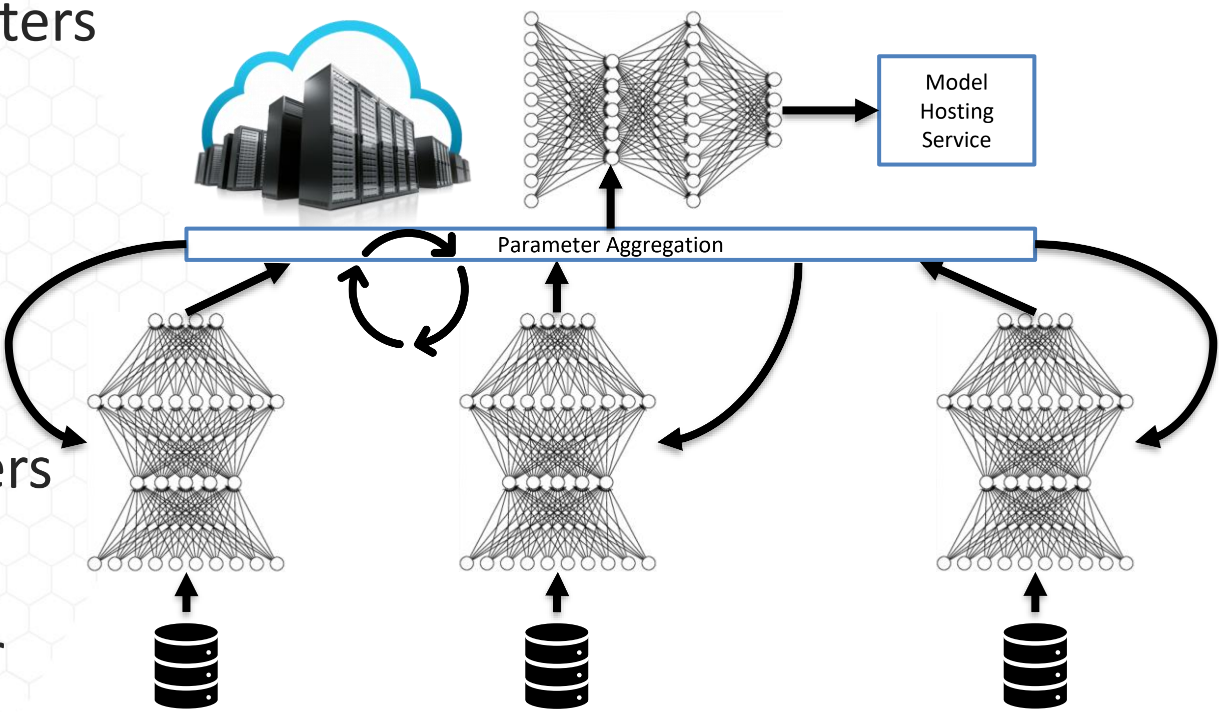

PRIVACY LEAKAGE IN FL SYSTEMS: SERVER, CLIENT, MODEL USER

Membership Inference Attacks:

Given training dataset D, and a model M trained on D, and a data point x.

Can an attacker determine if $x \in D$*?*

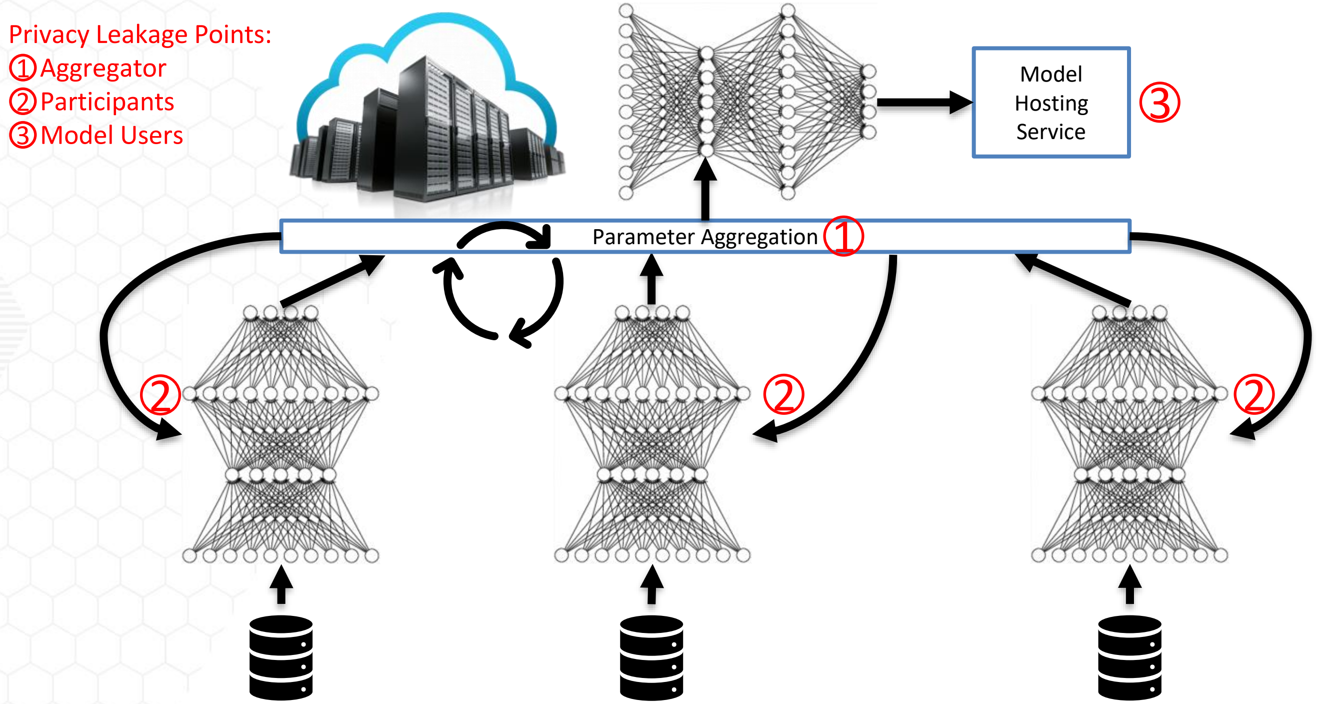

Attacker: Aggregator (Passive)

Dataset	Attack Accuracy[1]
Purchase History (100 class)	[illegible]
Texas Hospital Stays	[illegible]
CIFAR-100 (AlexNet)	[illegible]

[1] Nasr, M., R. Shokri, and A. Houmansadr. Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. *2019 IEEE Symposium on Security and Privacy (SP)*.

Definition

Differential Privacy [1]: A randomized mechanism K provides (ϵ, δ)-differential privacy if for any two neighboring databases D_1 and D_2 that differ in only a single entry and $\forall S \subseteq Range(K)$

$$\Pr(K(D_1) \in S) \leq e^{\epsilon} \cdot \Pr(K(D_2) \in S) + \delta$$

If $\delta = 0$, K is said to satisfy ϵ-differential privacy.

*****Limits the impact that any one instance can have on the mechanism output*****

[1] Dwork. Differential Privacy: A Survey of Results. 2008. International Conference on Theory and Applications of Models of Computation

Composition Property

Sequential Composition property[1]: Let $f_1, f_2, \dots, f_n$ be n algorithms such that for each $i \in [1, n]$, f_i satisfies (ϵ_i, δ_i)-differential privacy. Then,

Releasing the outputs of $f_1(D), f_2(D), \dots, f_n(D)$ satisfies $(\sum_{i=1}^{n} \epsilon_i, \sum_{i=1}^{n} \delta_i)$-DP.

*****Multiple passes on a dataset causes additive privacy loss in differential privacy*****

[1] Dwork et al. The algorithmic foundations of differential privacy. 2014. Foundations and Trends® in Theoretical Computer Science.

LOCAL DIFFERENTIAL PRIVACY

Definition

ϵ-LDP [1]: A randomized mechanism Ψ provides ϵ- local differential privacy where $\epsilon > 0$, if and only if for any inputs v_1, v_2 in universe $\mathcal{U}$ and $\forall y \in Range(\Psi)$, we have:

$$\Pr[\Psi(v_1) = y] \leq e^{\epsilon} \cdot \Pr[\Psi(v_2) = y]$$

*****Protects the raw value (input to Ψ) from privacy inference according to ϵ*****

[1] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. 2017. Collecting telemetry data privately. In Advances in Neural Information Processing Systems. 3571–3580.

α-CONDENSED LOCAL DIFFERENTIAL PRIVACY

Definition

α-CLDP [1]: A randomized mechanism Φ provides α- condensed local differential privacy where $\alpha > 0$, if and only if for any inputs v_1, v_2 in universe $\mathcal{U}$ and $\forall y \in Range(\Phi)$, we have:

$$\Pr[\Phi(v_1) = y] \leq e^{\alpha \cdot d(v_1, v_2)} \cdot \Pr[\Phi(v_2) = y]$$

*****Protects the raw value (input to Φ) from privacy inference according to α *and* $d(\cdot,\cdot)$*****

[1] M. Emre Gursoy, A. Tamersoy, S. Truex, W. Wei, and L. Liu. 2019. Secure and utility-aware data collection with condensed local differential privacy. IEEE Transactions on Dependable and Secure Computing (2019).

LDP-FED: FEDERATED LEARNING WITH LOCAL DIFFERENTIAL PRIVACY

Client Perspective

1. Participants initialize local DNN instance w/ θ_0 and each local LDP Module w/ individual privacy preferences.
2. Each participant P_i locally computes training gradients according to D_i.
3. Each P_i perturbs their gradients according to their local instance of the LDP Module.
4. The k-Client Selection Module accepts update from each P_i w/ probability $q = k/N$
5. Each participant waits to receive aggregated parameter updates from the parameter server and then updates its local models.
6. Each P_i proceeds to step 2 to start the next iteration.

Server Perspective

1. The parameter server generates initial model parameters θ_0 and sends to each participant.
2. The server waits to receive k parameter updates randomly selected by the k-Client Selection Module.
3. Once parameter updates are received, the Aggregation Module aggregates the updates, i.e., averages the gradient updates to determine new model parameters.
4. The parameter server updates model parameters and sends the updated values back to participants to update local DNN.

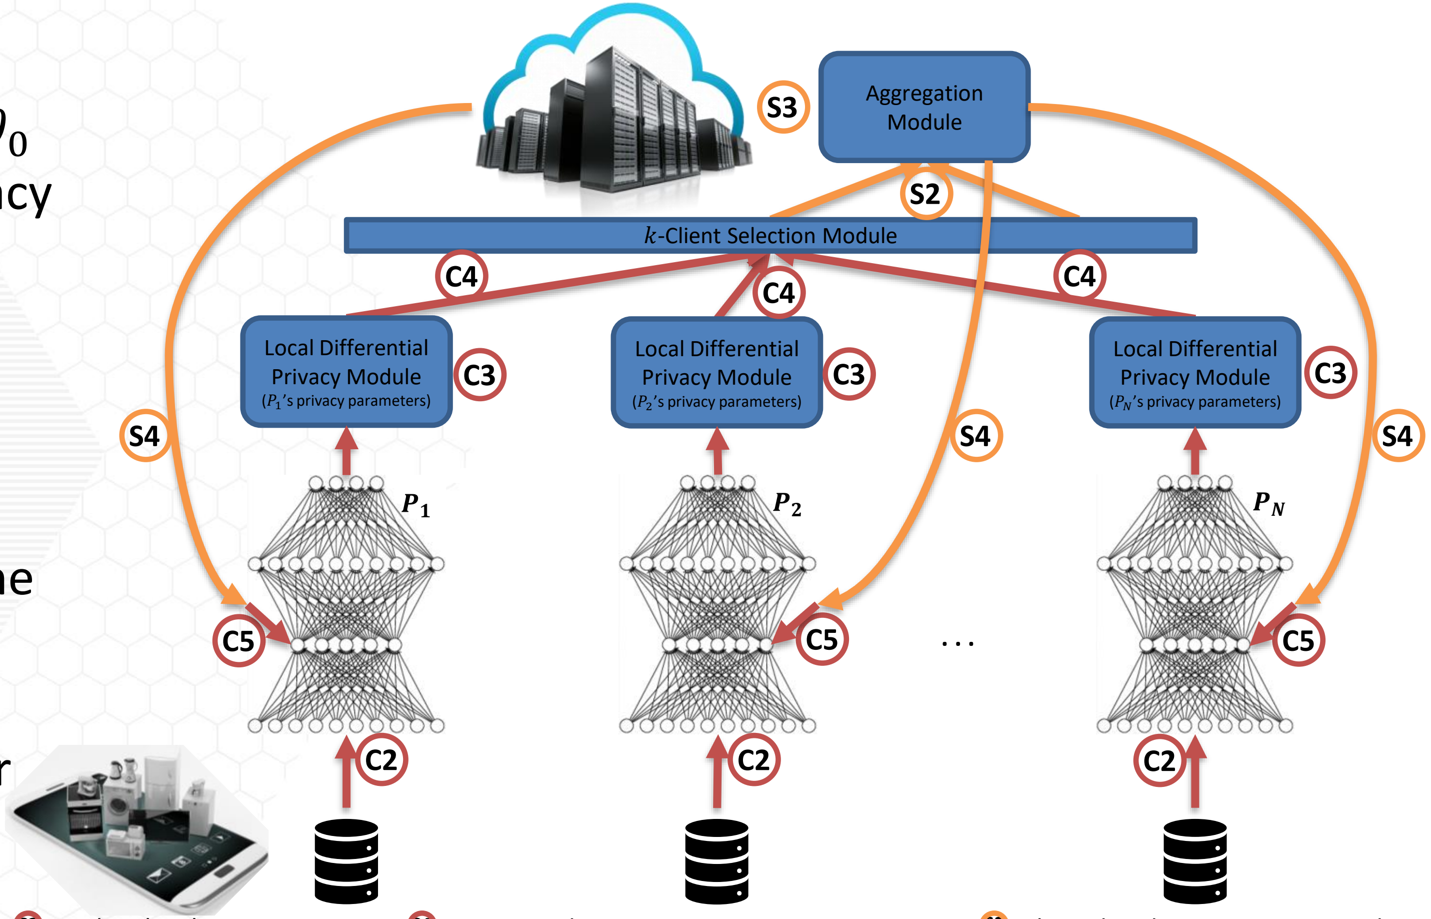

S3
Aggregation Module
S2
k-Client Selection Module
C4
C4
C4
Local Differential Privacy Module (P_1's privacy parameters)
C3
Local Differential Privacy Module (P_2's privacy parameters)
C3
Local Differential Privacy Module (P_N's privacy parameters)
C3
S4
S4
S4
P_1
P_2
P_N
C5
C5
C5
…
C2
C2
C2

C2 Conduct local training
C3 LDP Module to perturbs update
C4 Private updates sent to k-Client Selection Module
S2 Selection Module sends updates to Server
S3 Selected updates are aggregated
S4 New model parameters are sent back
C5 Local models are updated

LDP MODULE: PERSONALIZATION

- Individual participants locally define LDP-Module in LDP-Fed
  - Privacy guarantee, privacy mechanism parameters
- Privacy risk is not uniform:
  - Smaller datasets [1]
  - Minority group representation [2][3]
- Privacy requirements may not be uniform

Target Population	Attack Accuracy [3]
Aggregate	70.14%
Male Images	68.18%
Female Images	**76.85%**
White Race Images	62.77%
Racial Minority Images	**89.90%**

[1] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In 2017 IEEE Symposium on Security and Privacy (SP). IEEE, 3–18.

[2] Reza Shokri, Martin Strobel, and Yair Zick. Privacy risks of explaining machine learning models. arXiv preprint arXiv:1907.00164 (2019).

[3] Stacey Truex, Ling Liu, Mehmet Emre Gursoy, Wenqi Wei, and Lei Yu. Effects of Differential Privacy and Data Skewness on Membership Inference Vulnerability. arXiv preprint arXiv:1911.09777 (2019).

LDP MODULE: PRIVACY BUDGET ALLOCATION

- Privacy requirement: guarantee α-CLDP for each participant in FL training of DNN
- Must partition α into E small budgets! (one for each of the E total iterations) such that

$$\alpha = \sum_{i=0}^{E-1} \alpha_i$$

- Let $\theta_i =$ # of parameter updates to be uploaded to the parameter server at iteration i and α_i be the allocated portion of the overall privacy budget. We then set

$$\alpha_p = \frac{\alpha_i}{|\theta_i|}$$

- α_p is the privacy budget when applying α-CLDP to each parameter update in θ_i

LDP MODULE: PARAMETER SELECTION

- Basic implementation of α-CLDP in FL divides the budget by (1) number of iterations and (2) number of parameters in the model:

$$\alpha_p = \frac{\alpha}{qE|\theta|}$$

- Approach in α-CLDP-Fed is to reduce (2) to only upload a subset of the parameters θ_i at each iteration and therefore increase the budget α_p (and corresponding accuracy) for parameters which are uploaded
- In LDP-Fed: θ_i corresponds to 1 layer of the DNN with earlier iterations updating later layers and proceeding iterations moving backward through the network.
- Number of iterations and portion of the privacy budget allocated to an individual layer ℓ is directly proportionate to the size of that layer (with a minimum of 1 iteration)

LDP MODULE: CYCLES

- LDP-Fed cycles further control when different parameter updates are shared
- Each cycle is implemented in terms of iteration rounds

- Let $c' =$ number of cycles. One cycle is then $\frac{E}{c'}$ rounds.
- Rounds within each cycle are then allocated to individual layers in the same manner, with number of rounds allocated being proportional to layer size.

- In LDP-Fed, the default cycle value is set to 5.

LDP MODULE: EXTENDING LDP MECHANISM [1]

- Let θ_i be the parameters selected for upload by the LDP Module at iteration i
- For each parameter $p \in \theta_i$ the LDP Module then applies the appropriate LDP Mechanism; for α-CLDP-Fed…

Exponential Mechanism

Exponential Noise Mechanism[1]: Let $v \in \mathcal{U}$ be the raw user data, and let the Exponential Mechanism Φ_{EM} take as input v and output a perturbed value in $\mathcal{U}$, i.e.m $\Phi_{EM}\colon \mathcal{U} \rightarrow \mathcal{U}$. Then, Φ_{EM} that produces output y with the following probability satisfies α-CLDP.

$$\forall y \in \mathcal{U}\colon \Pr[\Phi_{EM}(v) = y] = \frac{e^{\frac{-\alpha \cdot d(v,y)}{2}}}{\sum_{z \in \mathcal{U}} e^{\frac{-\alpha \cdot d(v,z)}{2}}}$$

*****Add noise to each parameter value to achieve α-condensed local differential privacy*****

[1] M. Emre Gursoy, A. Tamersoy, S. Truex, W. Wei, and L. Liu. 2019. Secure and utility-aware data collection with condensed local differential privacy. IEEE Transactions on Dependable and Secure Computing (2019).

k-CLIENT SELECTION MODULE

- Conventional FL systems do not query every participant in every round
  - Efficiency
  - Availability (WiFi, power, etc.)
- Training in LDP-Fed: only $k \leq N$ participants' parameter updates selected per round
- Discarded updates do not introduce any privacy cost

Sampling Amplification

Allows for a tighter bound of $\alpha = \sum_{i=0}^{E-1} q \cdot \alpha_i$ where $q = \frac{k}{N} \leq 1$.

RESULTS: LDP-FED ACCURACY IN DL VS LDP-BASIC

$$\alpha = 1.0$$

- CLDP-Basic: below the random guess baseline of 10% $\Longrightarrow$ applying the privacy budget uniformly leads to untenable accuracy loss.
- LDP-Fed Single Layer approach significantly improves performance
- LDP-Fed's proportionate budget and iteration allocation further improves accuracy by an additional ~2%

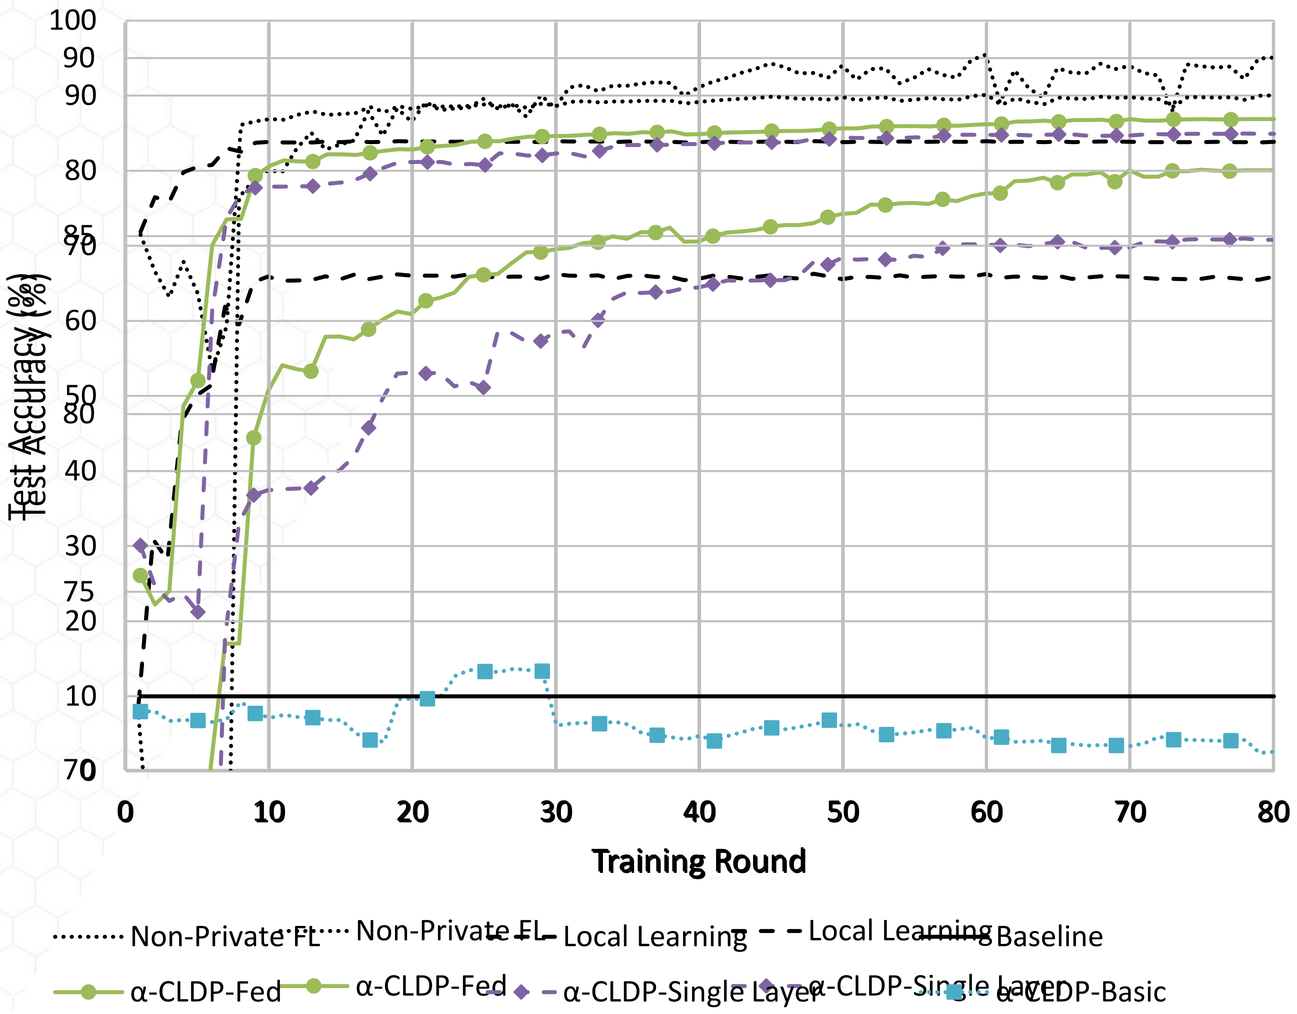

RESULTS: LDP-FED VS OPTIMIZER-BASED DP

$$\alpha = 1$$

- Adversarially equivalent ϵ computed from [1], $\delta = 10^{-5}$
- α-CLDP-Fed outperforms DPSGD by $\sim$6.8% and Hybrid-One by $\sim$3.5%
- Φ_{EM} in the LDP Module can be applied in parallel compared to cost of optimizer efficiency in DPSGD and Hybrid-One
- LDP-Fed requires no heavy cryptographic protocols

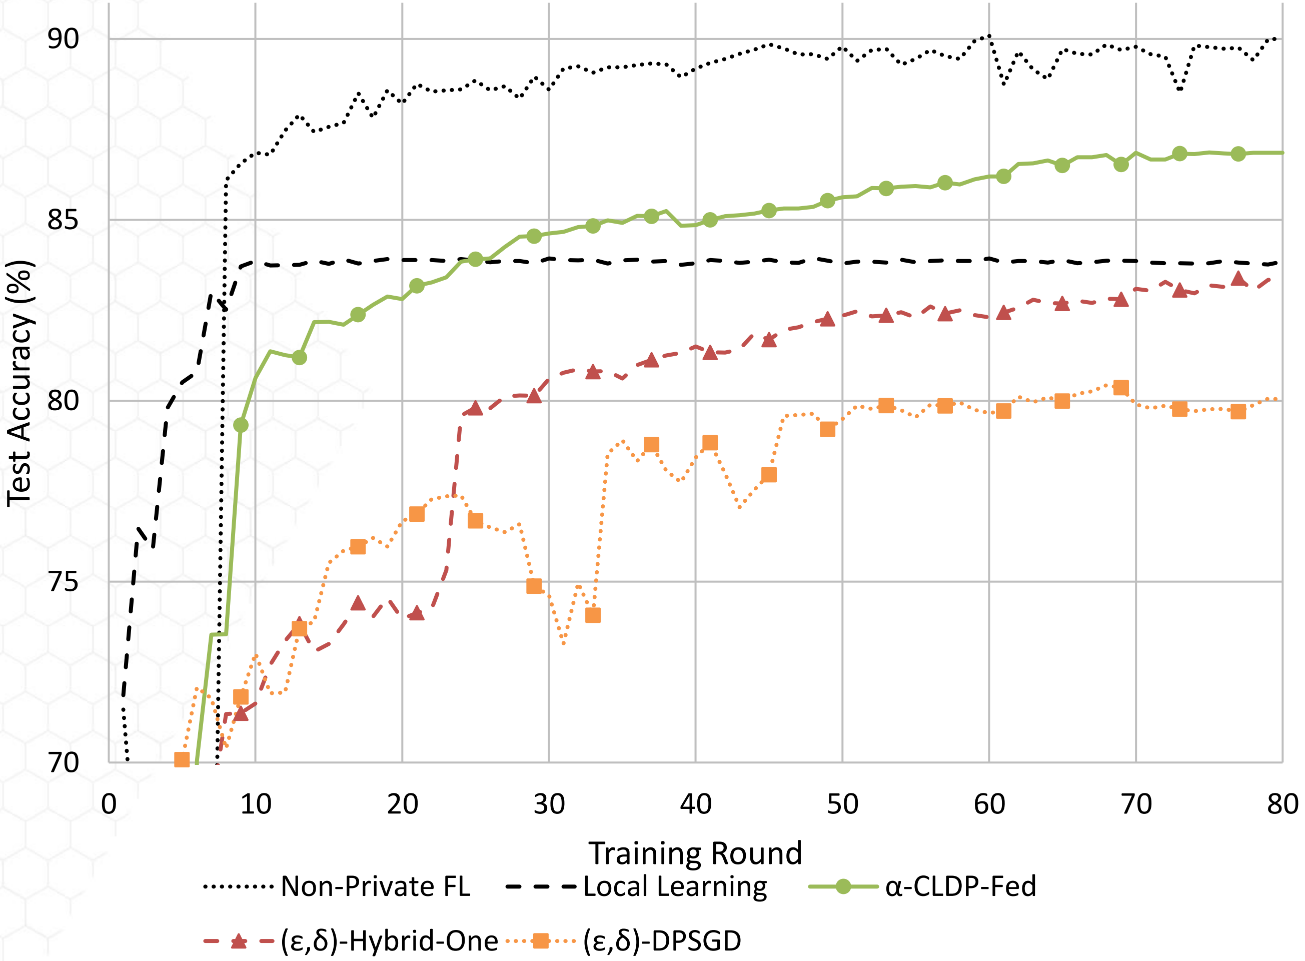

LDP-FED SYSTEM FEATURES

Privacy-Preserving Federated Learning Method	Efficient	Locally Defined Privacy Guarantee	Protection from Inference Attacks	Handles Complex Models
SMC [1]	🟥	🟥	🟨	🟩
ϵ-DP Paramater Sharing [2]	🟩	🟩	🟨	🟩
Local Optimizer [3]	🟨	🟩	🟩	🟥
Hybrid-One [4]	🟥	🟥	🟩	🟨
α-CLDP-Fed	🟩	🟩	🟩	🟩

[illegible]

- [illegible]
- [illegible] to each parameter in parallel
- Allows for adherence to local policies and compute restrictions

[1] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. ACM, 1175–1191

[2] Reza Shokri and Vitaly Shmatikov. 2015. Privacy-preserving deep learning. In Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. 1310–1321.

[3] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. 308–318

[4] Stacey Truex, Nathalie Baracaldo, Ali Anwar, Thomas Steinke, Heiko Ludwig, Rui Zhang, and Yi Zhou. 2019. A hybrid approach to privacy-preserving federated learning. In Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security. 1–11

CONCLUSION

- LDP-Fed: a novel FL approach with communication efficient LDP
  - An edge system for distributed and collaborative training with a large population of clients
  - Participants efficiently train complex models + formal privacy protection
  - Participants customize their LDP privacy budget locally
- The α-CLDP-Fed algorithm extends traditional LDP intended for single categorical values, to handle high dimensional, continuous, and large scale model parameter updates
- LDP-Fed parameter selection approach prevents LDP noise from overwhelming model updates → balancing utility, privacy trade-off
- Comparison of LDP-Fed with the state-of-the-art privacy-preserving FL approaches in both accuracy and system features.