On Differentially Private Federated Linear Contextual Bandits

Xingyu Zhou, Wayne State University

AI-EDGE Seminar @ OSU

Joint work with Sayak Ray Chowdhury, Microsoft Research, India

Mar 2, 2023

Cross-silo Federated Learning^[KMA+19] **A hospital example**













local model at each silo/agent



٠

Cross-silo Federated Learning A hospital example













local model at each silo/agent



٩

Cross-silo Federated Learning A hospital example













Cross-device FL

- Large no. of clients
- Limited resource
- e.g., clients are phones •

Cross-silo FL

- Small no. agents/silos
- More resource
 - e.g., silos are hospitals, banks, schools

updated local model at each silo/agent





In FL batch (supervised) learning:

- each silo has a static dataset
- offline training
- e.g., cancer diagnosis











time t

In FL **online** learning:

- each silo has a stream of data
- online training/decision, i.e., learn from interaction
- e.g., personalized medical care



٩





time t

In FL **online** learning:

- each silo has a stream of data
- online training/decision, i.e., learn from interaction
- e.g., personalized medical care











In FL **online** learning:

- each silo has a stream of data
- online training/decision, i.e., learn from interaction
- e.g., personalized medical care



٠

Cross-silo Federated Learning Linear contextual bandits (LCB)[APS11]



time t



In FL **LCBs** (online):

- unknown reward feedback *f* is a **linear** function $y_t = x_t^\top \theta^* + \eta_t$
- $x_t = \phi(c_t, a_t) \in \mathbb{R}^d$, ϕ is the feature map, c_t is context and a_t is the action
- θ^* is the unknown parameter
- η_t is zero-mean noise



Performance metric: group regret over M agents during T rounds $R_{M}(T) = \sum_{i=1}^{M} \sum_{t=1}^{T} \left[\max_{a} \phi_{i}(c_{t,i}, a)^{\top} \theta^{*} - \phi_{i}(c_{t,i}, a_{t,i})^{\top} \theta^{*} \right]$

Privacy in Cross-silo FL Though locally stored data, privacy risks still exist















Model parameters and updates can still reveal sensitive information e.g., model inversion attacks ^[FJR15, HZL19] or membership inference attacks ^[SSS+17]



Adversary could also be other silos .e., collude to infer users in another silo

Differentially Private Cross-silo FL Differential privacy^[DR14] — a rigorous privacy protection

Differential Privacy 101

Definition. If for any two neighboring datasets D and D', and any outcome E $\mathbb{P}(M(D) \in E) \leq e^{e}\mathbb{P}(M(D') \in E) + \delta$ Then, M satisfies (e, δ) -DP - DP means that outputs are "close" in

probability^[1] on two neighboring datasets

Key components:

- 1. What are the neighboring datasets?— the identity for protection
- 2. What are the outputs?
 - the view of adversary

Key properties:

- 1. Composition, privacy loss adds up
- 2. Post-processing, immune to further processing if data is not touched









¹¹ [1] A more general view is via *f*-divergence or Reny divergence between two distributions

Silo-level Local Differential Privacy (LDP) All communication from each silo is private

Silo-level LDP^[1]

Definition (informal). The full transcript of communication between any agent $i \in [M]$ and server are "close" in prob. on any two local neighboring datasets at agent i

Local neighboring datasets at agent *i*: *a*

sequence of T users that differs in only one user

- protect each user/patient
- different from standard DP for cross-device
- FL, where each client is protected

Outputs: full communication transcript

- communicated models/messages
- communication schedule, i.e., when
- communication happens













Private Federated LCB The state-of-the-art^[DP20]

[Dubey&Pentland '20]

Algorithm: federated LinUCB with Gaussian mechanism (tree-based)

Privacy: essentially the same as silo-level LDP

Regret: additional regret due to privacy is $\tilde{O}(\sqrt{MT/\epsilon})$

Conclusion: match the regret achieved by a "super" single agent

Fundamental Gaps

Privacy leakage

- The proposed algorithm fails to guarantee silo-level LDP
- A simple attack can reveal sensitive information of users

Incorrect regret

- The claimed privacy cost is mis-calculated
- The correct one is $\tilde{O}(M^{3/4}\sqrt{T/\epsilon})$
- Hence, no longer match the "lower bound"



[DP20] Differentially-private federated linear bandits. NeurIPS 2020



Contribution

Main Results



- Identify the privacy and regret gaps in the state-of-the-art 1.
- 2. Propose a generic federated algorithm with flexible privacy protocols
- Achieve the correct regret bound under silo-level LDP, i.e., the privacy cost is $\tilde{O}(M^{3/4}\sqrt{T/\epsilon})$ 3.
- Shave the additional $M^{1/4}$ factor under shuffle differential privacy (SDP) still a weak trust DP model 4.

















The communication **schedule** is not fully private Observe when sync happens, other silos can infer the user in another silo















Communication schedule for silos in SOTA 📩

$$\exists i \in [M], \quad f(X_i, Z) > 0$$

• Z - all previous sync data among all silos

- X_i **new non-private** local data at silo *i* since recent sync
- f sync function, shared among all silos















Communication schedule for silos in SOTA 📩

$$\exists i \in [M], \quad f(X_i, Z) > 0$$

- Z all previous sync data among all silos
- X_i **new non-private** local data at silo *i* since recent sync
- f sync function, shared among all silos















Communication schedule for silos in SOTA

$$\exists i \in [M], \quad f(X_i, Z) > 0$$

- Z all previous sync data among all silos
- X_i **new non-private** local data at silo *i* since recent sync
- sync function, shared among all silos



Malicious silo can take advantage of this to infer user's sensitive data in another silo

Privacy Gap in SOTA A simple toy-example attack



time t = 1

Communication schedule for silos in SOTA

 $\exists i \in [M], \quad f(X_i, Z) > 0$

- Z all previous sync data among all silos
- X_i **new non-private** local data at silo *i* since recent sync
- sync function, shared among all silos

$$_{\rm e}, Z = 0) > 0$$







Regret Gap in SOTA Miscalculated total privacy noise













Larger total privacy noise implies larger regret

Ignore the privacy issue, the total amount of privacy noise in SOTA needs to be $\sigma_{\text{total}}^2 = M \sigma^2$, i.e., M factor of its current one (recall M is the no. silos)



Current conclusion in SOTA becomes ungrounded

After the correction of M factor, the regret due to privacy changes from

 $\tilde{O}(\sqrt{MT/\epsilon})$ (match the "lower bound" of a single agent)

to

 $\tilde{O}(M^{3/4}\sqrt{T/\epsilon})$ (has a gap of $M^{1/4}$ compared to "lower bound")





Motivating Questions

- 2. How to correct the regret bound while preserving the privacy?
- 3. How to close the gap compared to the "lower bound"? (\bigcirc need a way to get rid of M factor)
- 4. If possible, can we achieve all of them via a generic method? (? a template algorithm with a template proof is preferred)



1. How to address the privacy leakage? (a fixed communication schedule may work, i.e., does not depend on user's non-private data)

Our Approach

A Generic Algorithm

Private-FedLinUCB

Private-FedLinUCB (fixed batch sync of LinUCB with privacy) **Parameters**: batch size *B*, privacy protocol P = (R, S, A)Initialize: $\forall i, W_i = 0, U_i = 0; \ \widetilde{W}_{syn} = 0, \ \widetilde{U}_{syn} = 0$ for t = 1, ..., T do for each agent i = 1, ..., M do $V_{t,i} = \lambda I + \widetilde{W}_{syn} + W_i, U_{t,i} = \widetilde{U}_{syn} + U_i$ Estimate: $\hat{\theta}_{t,i} = V_{t,i}^{-1}U_{t,i}$ UCB: $a_{t,i} = \arg \max_{a} \phi(c_{t,i}, a)^{\mathsf{T}} \hat{\theta}_t + \beta_t \| \phi(c_{t,i}, a) \|_{V_{\mathsf{T}}^{-1}}$ Observe reward $y_{t,i}$; set $x_{t,i} = \phi(c_{t,i}, a_{t,i})$ Update local data: $W_i = W_i + x_{t,i} x_{t,i}^{\top}$, $U_i = U_i + x_{t,i} y_{t,i}$ if $t \mod B = 0$ then $\widetilde{W}_{syn} = P(\{W_i\}_{i \in [M]}), \ \widetilde{U}_{syn} = P(\{U_i\}_{i \in [M]}) \checkmark$ Receive $\widetilde{W}_{\text{SVN}}$, $\widetilde{U}_{\text{SVN}}$ from server Reset $W_i = 0, U_i = 0$

Single agent LinUCB^[APS11] 101

For t = 1, ..., T: **1. Estimate** θ^* : $\hat{\theta}_t = V_t^{-1}U_t$, $(V_t = \lambda I + \sum_{s=1}^{t-1} x_s x_s^{\top}$ ("covariance"), $U_t = \sum_{s=1}^{t-1} x_s y_s$ ("bias")) **2.** UCB: $a_t = \arg \max_a \phi(c_t, a)^{\mathsf{T}} \hat{\theta}_t + \beta_t \| \phi(c_t, a) \|_{V_t^{-1}}$ $(x_t = \phi(c_t, a_t), \beta_t - \text{chosen via confidence bound})$

 W_i – sum of local covariance matrices at agent i

 U_i – sum of local bias vectors at agent i

 W_{syn} — private sync covariance matrices among all agents

 U_{syn} — private sync bias vectors among all agents



P = (R, S, A), a template protocol for **summation** (will discuss it soon) R – local randomzier at agent side (on W_i , U_i) S — shuffler or identity mapping, between agents, server -->

A — analyzer at server side



A Generic Algorithm

Private-FedLinUCB

Private-FedLinUCB (fixed batch sync of LinUCB with privacy) **Parameters**: batch size *B*, privacy protocol P = (R, S, A)Initialize: $\forall i, W_i = 0, U_i = 0; \ \widetilde{W}_{syn} = 0, \ \widetilde{U}_{syn} = 0$ for t = 1, ..., T do for each agent i = 1, ..., M do $V_{t,i} = \lambda I + \widetilde{W}_{syn} + W_i, U_{t,i} = \widetilde{U}_{syn} + U_i$ Estimate: $\hat{\theta}_{t,i} = V_{t,i}^{-1}U_{t,i}$ UCB: $a_{t,i} = \arg \max_{a} \phi(c_{t,i}, a)^{\mathsf{T}} \hat{\theta}_t + \beta_t \| \phi(c_{t,i}, a) \|_{V^{\mathsf{T}}}$ Observe reward $y_{t,i}$; set $x_{t,i} = \phi(c_{t,i}, a_{t,i})$ Update local data: $W_i = W_i + x_{t,i} x_{t,i}^{\top}$, $U_i = U_i + x_{t,i} y_{t,i}$ if $t \mod B = 0$ then $W_{syn} = P(\{W_i\}_{i \in [M]}), \ U_{syn} = P(\{U_i\}_{i \in [M]})$ Receive \widetilde{W}_{syn} , \widetilde{U}_{syn} from server Reset $W_i = 0, U_i = 0$

Single agent LinUCB^[APS11] 101

For
$$t = 1, ..., T$$
:
1. Estimate $\theta^*: \hat{\theta}_t = V_t^{-1} U_t$,
 $(V_t = \lambda I + \sum_{s=1}^{t-1} x_s x_s^{\top} (\text{``covariance''}), U_t = \sum_{s=1}^{t-1} x_s y_s$
2. UCB: $a_t = \arg \max_a \phi(c_t, a)^{\top} \hat{\theta}_t + \beta_t \| \phi(c_t, a) (x_t = \phi(c_t, a_t), \beta_t) - c$

Remark: fixed vs. adaptive schedule

- It now suffices to privatize each sent messages for silo-level LDP guarantee
 - $-R(W_i)$ and $R(U_i)$ is private at each sync round $k \in [T/B]$
 - without worrying privacy leakage via schedule
 - -B needs to balance between comm. cost, regret, and privacy
- The problem in SOTA is: schedule depends on **non-private data** (i.e., W_i)
 - how about privatizing it first and then be adaptive?
 - we show that it will lead to fundamental challenge in regret analysis

P = (R, S, A), a template protocol for **summation** (will discuss it soon) R — local randomzier at agent side (on W_i , U_i)

- S shuffler or identity mapping, between agents, server
- A analyzer at server side





A Generic Priv. Protocol Distributed Tree-based alg.

P = (R,S,A), privacy protocol (distributed version of Tree-based algorithm)

Differential Privacy 201

- 1. Gaussian mechanism for private sum of l_2 bounded vectors i.e., \tilde{s} is the private sum of $\sum \gamma_i$ under (ϵ, δ) -DP $\widetilde{s} = \sum_{i=1}^{n} \gamma_{i} + \mathcal{N}(0, \sigma^{2}I), \sigma^{2} \approx \frac{L^{2} \log(1/\delta)}{\sigma^{2}}$ **Intuition**: change one data, the sum changes in l_2 , bounded by L
- 2. Continual private sum (essential for private online learning)

i.e., a stream of data $\gamma_1, \ldots, \gamma_K$, compute \widetilde{s}_k – priv. sum of $\sum \gamma_s$

<u>Simple Approach I:</u> add noise ($\approx 1/\epsilon^2$) to each γ_s

- $-(\epsilon, \delta)$ -DP (by post-processing)
- total noise is K/ϵ^2 (!)

<u>Simple Approach II:</u> add noise ($\approx 1/\epsilon^2$) to each prefix sum

— total noise is $1/\epsilon^2$ for all k

$$- \approx (\sqrt{K\epsilon, \delta'})$$
-DP (by composition of DP)

- i.e., for (ϵ, δ) -DP, the total noise needs to be K/ϵ^2 (!)



A Generic Priv. Protocol Distributed Tree-based alg.

P = (R, S, A), privacy protocol (distributed version of Tree-based algorithm)



A Generic Priv. Protocol Distributed Tree-based alg.

P = (R,S,A), privacy protocol (distributed version of Tree-based algorithm)

Differential Privacy 201



A Generic Priv. Protocol

Distributed Tree-based alg.

P = (R, S, A), privacy protocol (distributed version of Tree-based algorithm)

Procedure: Local Randomizer R at each agent $i \in [M]$

for each sync $k = 1, \dots, K$ do

Express k in binary form: $k = \sum Bin_j(k) \cdot 2^j$

Find index of first one $i_k = \min\{j : Bin_j(k) = 1\}$ Compute non-private p-sum: $\alpha_{i_k} = \sum \alpha_j + \gamma_k$

Output noisy p-sum $\widetilde{\alpha}_{i_{k},i} = \alpha_{i_{k}} + \mathcal{N}(0,\sigma^{2}I)$

Procedure: Shuffler *S* (could be empty or identity mapping)

Procedure: Analyzer *A* at server

for each sync $k = 1, \dots, K$ do

Express k in binary form and find index of first one i_k

Add noisy p-sums from all agents $\widetilde{\alpha}_{i_k} = \sum \widetilde{\alpha}_{i_k,i_k}$ $i \in [M]$

Output total sum: $\widetilde{s}_k = \sum_{i=1}^{k} \widetilde{\alpha}_i$ *j*:Bin(k)=1

Differential Privacy 201





Putting them together

• Each agent runs two privacy protocol – sum of **covariance** matrices (i.e., W_i) and sum of **bias** vectors (i.e., U_i)

• The datapoint γ_k is a **batch** of data — total matrices or vectors during the kth batch

• The sensitivity does not scale with respect to *B*



Procedure: Analyzer *A* at server

for each sync $k = 1, \ldots, K$ do

Express k in binary form and find index of first one i_k Add noisy p-sums from all agents $\widetilde{\alpha}_{i_{\nu}} = \sum \widetilde{\alpha}_{i_{\nu},i_{\nu}}$

 $i \in [M]$

Output to

tal sum:
$$\widetilde{s}_k = \sum_{j: \mathsf{Bin}(\mathsf{k})=1} \widetilde{\alpha}_j$$

 $x_t x_t^{\mathsf{T}}$

 $x_t y_t$

 \sum



Algorithm in action Illustration $\gamma_6^{\text{bias}} =$ $x_t y_t$





How about summing over time at

Private sum across both time and agents $\sum_{i=1}^{M} \sum_{j=1}^{6B} x_{t,i} y_{t,i}$ i=1 t=1





An alternative protocol Palt

time
$$t = 6B$$



٩

An alternative protocol Palt $\frac{6B}{\sum}$ $\gamma_6^{\text{bias}} =$ $x_t y_t$





- However, for shuffle DP, things are different
 - our protocol manages to close the gap
 - $-P_{\text{alt}}$ fails to close the gap (more on this later...





Theoretical Results



Federated LCBs under Silo-level LDP Fix the issues in SOTA

Theorem 1 (Performance under silo-level LDP, informal)

Let batch size $B = \sqrt{T/M}$, privacy noise in P be $\sigma^2 = 8\kappa \cdot \frac{\log(2/\delta) + \epsilon}{\epsilon^2}$ with $\kappa = 1 + \log(T/B)$. Then, Private-FedLinUCB enjoys

- 1. **Privacy** (ϵ, δ) -silo-level LDP for any $\epsilon > 0, \ \delta \in (0,1)$
- 2. **Regret** $-R_M(T) = \text{non-private regret} + \sqrt{T} \frac{(Md)^{3/4} \log^{1/4} (1/\delta)}{\sqrt{\epsilon}}$
- 3. Communication $-\sqrt{MT}$ rounds of sync between agents and server

Remark: comparisons with related work

1. Compared with SOTA^[DP20]

- privacy: we fix the privacy leakage, thanks to the fixed-batch schedule and tree-based algorithm

- 2. Compared with "super" single agent under central DP[SS18]
 - our regret is $M^{1/4}$ factor worse than this "lower bound"

- regret: we establish the correct privacy cost, i.e., the additional regret due to privacy now scales with $M^{3/4}$ (instead of \sqrt{M})

- communication: communication is worse than SOTA (\sqrt{Tvs} . log T) due to fixed-batch comm. But, note that there exists privacy leakage

Federated LCBs under SDP Match the "lower bound"

Differential Privacy 501

1. What is shuffle DP (SDP)?

- formally defined in [CSUZZ19]
- -P = (R, S, A), "the output of shuffler is private"
- (change any d_i , the outputs are "close")



2. How to achieve it?

- one way is via LDP amplification, e.g., [FMT20]
- shuffle *n* LDP outputs (each ϵ_0 -DP), then it is $\approx \epsilon_0 / \sqrt{n}$ SDP
- "reduce the privacy loss by a factor of $1/\sqrt{n}$
- (intuition: hiding among clones)



Federated LCBs under SDP Match the "lower bound"



How about adding shuffler between agents and server? $\overline{1}/\sqrt{M}$

Good news: this amplification can close the gap \checkmark

Bad news: one cannot directly use existing results

- they only amplify LDP (R oper. on single data)
- in our case, R oper. on multiple datapoints
- (this leads to key difference in the analysis)

Clones are harder to create due to multiple local points

A new amplification lemma is derived 🗸

- tailored for Gaussian DP mecha.
- avoid group privacy
- control the blow up in δ

Differential Privacy 501

- 1. What is shuffle DP (SDP)?
 - formally defined in [CSUZZ19]
 - -P = (R, S, A), "the output of shuffler is private"
 - (change any d_i , the outputs are "close")



2. How to achieve it?

- one way is via LDP amplification, e.g., [FMT20]
- shuffle *n* LDP outputs (each ϵ_0 -DP), then it is $\approx \epsilon_0 / \sqrt{n}$ SDP
- "reduce the privacy loss by a factor of $1/\sqrt{n}$.
- (intuition: hiding among clones)





Federated LCBs under SDP Match the "lower bound"

Theorem 2 (Performance under SDP, informal)

3. Communication $-\sqrt{MT}$ rounds of sync between agents and server

Match the "lower bound"

Privacy cost is on the order of \sqrt{MT}

 $-1/\sqrt{M}$ is the standard term

How to improve the privacy guarantees?



Federated LCBs under SDP Leverage vector-sum protocol

Differential Privacy 502

1. How to achieve SDP?

- instead of using amplification lemma
- one can use specific shuffle protocol

 $-P_{\text{vec}} = (R_{\text{vec}}, S, A_{\text{vec}})^{[CJMP21]}$ is one example



2. Performance of P_{vec}

- it guarantees SDP for all $\epsilon \in (0,15), \delta \in (0,1/2)$

- the injected noise is $\frac{L^2}{\epsilon^2} \log^2(d/\delta)$ per entry (indep. of n)

(Essentially, it simulates central model without a trusted server)



Federated LCBs under SDP Leverage vector-sum protocol



The norm of p-sum could be linear with T

- sum of M p-sums with P_{vec} (i.e., n = M)
- each data point has a large norm

View *n* in P_{Vec} as data points across agents

- -e.g., for k = 6
- each p-sum has 2B points
- $-n = M \cdot 2B$ with each norm bounded
- each sync incurs only $1/\epsilon^2$ noise \checkmark

Differential Privacy 502



2. Performance of P_{vec}

1. How to achieve SDP?

- instead of using amplification lemma

- one can use specific shuffle protocol

- it guarantees SDP for all $\epsilon \in (0,15), \delta \in (0,1/2)$

- the injected noise is $\frac{2}{c^2} \log^2(d/\delta)$ per entry (indep. of n)

(Essentially, it simulates central model without a trusted server)



Algorithm in action With P_{vec}

time t = 6B

(the 6-th communication)



42

Federated LCBs under SDP Improved privacy via P_{vec}

Theorem 3 (Performance under SDP via *P***vec**, informal)

Let batch size $B = \sqrt{T/M}$ and $\kappa = 1 + \log(T/B)$. Combine P_{Vec} with our privacy protocol. Then, Private-FedLinUCB enjoys

1. **Privacy** – (ϵ, δ) -SDP for any $\epsilon \in (0, 60), \delta \in (0, 1), \blacktriangleleft$

2. **Regret** $-R_M(T) = \text{non-private regret} + \sqrt{MT} \frac{d^{3/4} \log^{3/4} (M\kappa/\delta)}{\sqrt{\epsilon}}$

3. Communication $-\sqrt{MT}$ rounds of sync between agents and server







Analysis

A Generic Analysis "One-line" proof for regret



Proposition 1 (Generic regret bound under PNC, informal)

Suppose that the privacy protocol satisfies PNC with parameter σ_{tot}^2 , then Private-FedLinUCB enjoys the following regret with high probability

$$R_{M}(T) = \tilde{O}\left(dMB + d\sqrt{MT} + \sqrt{\sigma_{tot}MT}d^{3/4}\right)$$

Cost due to batching Standard regret Cost due to privacy









Importance of p-sum Why *P*_{alt} fails for SDP

Prop. 1. Regret due to privacy: $\sqrt{\sigma_{tot}MT}$

Regret under SDP: $\tilde{O}(\sqrt{MT/\epsilon})$

•





Simulations

Discussions

Q1: Can we further reduce comm. cost to $\log T$

Q3: What if users even do not trust each local agent?

Then, it might need adaptive update based on determinant condition. Challenges exist in private case

Q2: Silo-level LDP/ SDP vs. other privacy notions in contextual bandits? Q4: What if users participate multiple times ? (within one silo or across silos)

We give a comprehensive discussions on difference and connections Q5: How to balance between privacy and algorithm complexity?

It turns out that a simple tweak of our algorithm can handle this situation

Good question. We are working on it right now

Q6: Can we generalize it to federated RL

One can use composition or group privacy to handle. Or directly analyze the total sensitivity

Yes, at least for RL with linear function approximation

One last thing...

Recent Research... Many thanks to all my collaborators

Private MAB

- "MAB under local DP with tight lower bound" [RZLS20, arxiv]
- "the state-of-the-art of private MAB for all three DP models" [CZ*23, ICLR23]
- "private and robust MAB" [WZ*TW23, submitted]

Private Contextual Bandits

- "linear contextual bandits under shuffle model" [CZ*22, ICML22]
- "federated LCBs under both silo-level LDP and SDP [ZC, arixv, submitted]
- "kernel bandits under local model" [ZT21, AAAI21]
- "private linear bandits with distributed feedback" [LZJ22, WiOpt22, Best Student Paper]
- "private distributed kernel bandits" [LZJ23, Sigmetrics23]
- Private RL
 - "A comprehensive study of tabular RL under both central and local DP models" [CZ*22, AAAI22, oral]
 - "The first study of private RL with linear function approximation" [Z22, Sigmetrics22]
 - "Study of private LQR" [CZ*S21, ISIT21]

Many interesting open problems in this area... Collaborations are welcome 🎉

Reference

- [KMA+19] Peter Kairouz, H. Brendan McMahan, Brendan Avent, et al. Advances and open problems in federated learning. arXiv preprint: 1912.04977, 2019
- [APS11] Abbasi-Yadkori, Yasin, Dávid Pál, and Csaba Szepesvári. "Improved algorithms for linear stochastic bandits." NeurIPS, 2011.
- ACM SIGSAC Conference on Computer and Communications Security, pages 1322–1333, 2015.
- Conference, pages 148–162, 2019
- 22 (pp. 3-18). IEEE.
- [LR21] Private federated learning without a trusted server: Optimal algorithms for convex losses. arXiv preprint arXiv:2106.09779, 2021
- [LHW+22] On privacy and personalization in cross-silo federated learning. arXiv preprint arXiv:2206.07902, 2022
- International Publishing.
- Annual Symposium on Foundations of Computer Science (FOCS), pp. 954-964. IEEE, 2022
- [CJMP21] Cheu, Albert, Matthew Joseph, Jieming Mao, and Binghui Peng. "Shuffle private stochastic convex optimization." arXiv preprint arXiv:2106.09805 (2021).
- [RZLS20] Ren, Wenbo, Xingyu Zhou, Jia Liu, and Ness B. Shroff. "Multi-armed bandits with local differential privacy." arXiv preprint arXiv:2007.03121 (2020).
- [CZ*23] Chowdhury, Sayak Ray, and Xingyu Zhou. "Distributed Differential Privacy in Multi-Armed Bandits." ICLR23.
- [WZ*TW] Wu, Yulian, Xingyu Zhou, Youming Tao, and Di Wang. "On Private and Robust Bandits." arXiv preprint arXiv:2302.02526 (2023).
- [CZ*22] Chowdhury, Sayak Ray, and Xingyu Zhou. "Shuffle private linear contextual bandits." ICML22.
- [ZC23] Xingyu Zhou, and Chowdhury, Sayak Ray "On Differentially Private Federated Linear Contextual Bandits"
- 11152-11159.2021.

• [FJR15] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In Proceedings of the 22nd

• [HZL19] Zecheng He, Tianwei Zhang, and Ruby B Lee. Model inversion attacks against collaborative inference. In Proceedings of the 35th Annual Computer Security Applications

• [SSS+17] Shokri R, Stronati M, Song C, Shmatikov V. Membership inference attacks against machine learning models. In 2017 IEEE symposium on security and privacy (SP) 2017 May

• [DR14] Dwork, Cynthia, and Aaron Roth. "The algorithmic foundations of differential privacy." Foundations and Trends® in Theoretical Computer Science 9.3–4 (2014): 211-407.

• [CSS11] Chan, T-H. Hubert, Elaine Shi, and Dawn Song. "Private and continual release of statistics." ACM Transactions on Information and System Security (TISSEC) 14.3 (2011): 1-24. • [CSUZZ19] Cheu, A., Smith, A., Ullman, J., Zeber, D., & Zhilyaev, M. (2019). Distributed differential privacy via shuffling. In Advances in Cryptology-EUROCRYPT 2019: 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part I 38 (pp. 375-403). Springer

• [FMT20] Feldman, Vitaly, Audra McMillan, and Kunal Talwar. "Hiding among the clones: A simple and nearly optimal analysis of privacy amplification by shuffling." In 2021 IEEE 62nd

• [ZT21] Zhou, Xingyu, and Jian Tan. "Local differential privacy for bayesian optimization." In Proceedings of the AAAI Conference on Artificial Intelligence, vol. 35, no. 12, pp.

Reference

- Optimization in Mobile, Ad hoc, and Wireless Networks (WiOpt), pp. 41-48. IEEE, 2022.
- [LZJ23] Li, Fengjiao, Xingyu Zhou, and Bo Ji. "(Private) Kernelized Bandits with Distributed Biased Feedback." arXiv preprint arXiv:2301.12061 (2023).
- Intelligence, vol. 36, no. 6, pp. 6375-6383. 2022.
- no. 1 (2022): 1-27.
- *Theory (ISIT*), pp. 485-490. IEEE, 2021.

• [LZJ22] Li, Fengjiao, Xingyu Zhou, and Bo Ji. "Differentially private linear bandits with partial distributed feedback." In 2022 20th International Symposium on Modeling and

• [CZ*22] Chowdhury, Sayak Ray, and Xingyu Zhou. "Differentially private regret minimization in episodic markov decision processes." In Proceedings of the AAAI Conference on Artificial

• [Z22] Zhou, Xingyu. "Differentially private reinforcement learning with linear function approximation." Proceedings of the ACM on Measurement and Analysis of Computing Systems 6,

• [CZ*] Chowdhury, Sayak Ray, Xingyu Zhou, and Ness Shroff. "Adaptive control of differentially private linear quadratic systems." In 2021 IEEE International Symposium on Information

Thank you!