Oracle Efficient Private Non-Convex Optimization

Seth Neel
University of Pennsylvania

Aaron Roth
University of Pennsylvania

Giuseppe Vietri
University of Minnesota

Steven Wu
University of Minnesota

First Setting.

- Consider following general optimization problem defined in terms of dataset $D \in X^n$.

$$\min_{w \in \mathscr{W}} L(D, w)$$

- $\mathscr{W}$ is a **discrete** and bounded set. For this talk assume that $||\mathscr{W}||_\infty = 1$
- The loss function $L : X^n \times \mathscr{W} \to [0,1]$ is **bounded**.

Can we solve solve this problem with differential privacy?

First Contribution

- We propose an Objective Perturbation Algorithm that:
  - Satisfies (ε, δ)-differential privacy.
  - W.h.p finds an answer $\widehat{w} \in \mathscr{W}$ with error bounded by

$$\left| L(D, \widehat{w}) - L(D, w^*) \right| \leq \frac{14\sqrt{2(d+1)\ln(2\beta)}\sqrt{\ln(1/\delta)}}{n\tau\epsilon}.$$

The Normalization Trick

- Let $\|w\|_2 \leq D$, For all $w \in \mathscr{W}$. The normalization function is:

$$\pi(w) = \left(w_1, \ldots, w_d, D\sqrt{1 - \|w\|_2/D}\right)\frac{1}{D}$$

- Key idea: Sample random vector η, and Augment objective:

$$\min_{w \in \mathscr{W}} L(D, w) - \langle \eta, \pi(w) \rangle$$

- Solve using non-private optimization oracle.

Stability of the Objective

- Let $w, w' \in \mathscr{W}$. Then

$$\|\pi(w) - \pi(w')\|_2^2 \geq \frac{1}{D^2}\|w - w'\|_2^2$$

- The normalization trick provides stability.

Objective Perturbation.

Algorithm 1: ObjPertDiscrete

Input: Projection function π, $\mathcal{D} = \{l_i\}$, optimization oracle

$n \leftarrow |\mathcal{D}|$;

$\sigma \leftarrow \frac{7LD^2\sqrt{\ln 1/\delta}}{\tau\epsilon}$;

Draw i.i.d random vector $\eta \sim \mathcal{N}\left(0, \sigma^2\right)^{d+1}$;

$$w \in \arg\min_{w \in \mathcal{W}_\tau} \left(\frac{L\left(\mathcal{D}, w\right) - \left\langle \eta, \pi\left(w\right) \right\rangle}{n} \right)$$

Output: w

Second Setting

- Consider following general optimization problem defined in terms of dataset $D \in X^n$.

$$\min_{w \in \mathscr{W}} L(D, w)$$

- $\mathscr{W}$ is a **general** decision space. For this talk assume that $||\mathscr{W}||_\infty = 1$
- The loss function $L : X^n \times \mathscr{W} \to [0,1]$ is G-**Lipschitz**, **convex** and **bounded**.

Can we solve solve this problem with differential privacy?

Second Contribution

- We propose an Objective Perturbation Algorithm that:
  - Satisfies (ε, δ)-differential privacy.
  - W.h.p finds an answer $\widehat{w} \in \mathscr{W}$ with error bounded by

$$\left| L(D, \widehat{w}) - L(D, w^*) \right| \leq \frac{d^{5/4} G \sqrt{D_2 \log(1/\beta)}}{\sqrt{n\epsilon}}$$

Standard Differential Privacy (DP)

[Dwork et al., 2006]

Two datasets

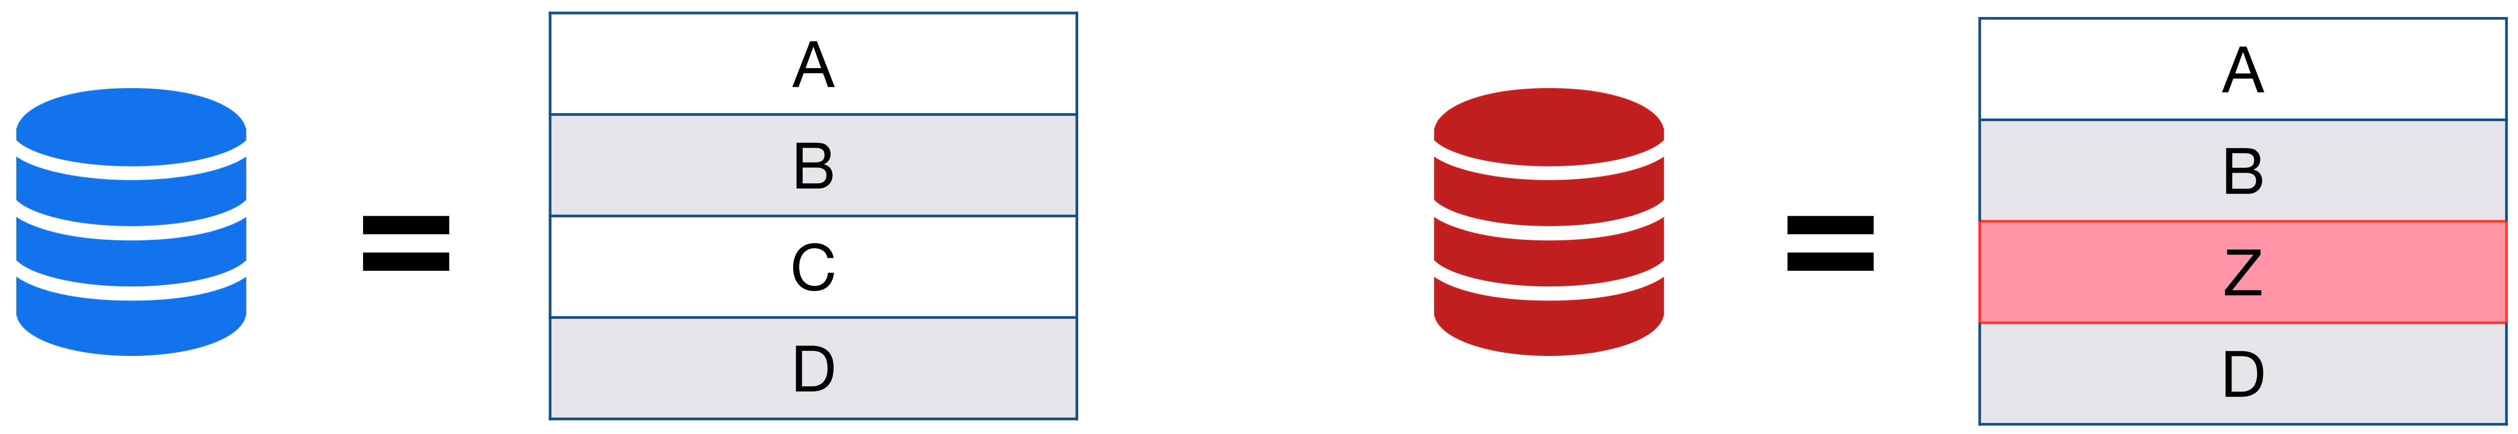

are **neighbors** if the are different on only one row.

Definition: Mechanism M satisfies ε-differential privacy if, for all neighboring datasets and for all $r \in \text{range}(M)$

$\Pr[M($ 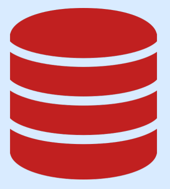 $) = r] \leq e^{\varepsilon} \Pr[M($  $) = r]$

Experiments.

Experiment

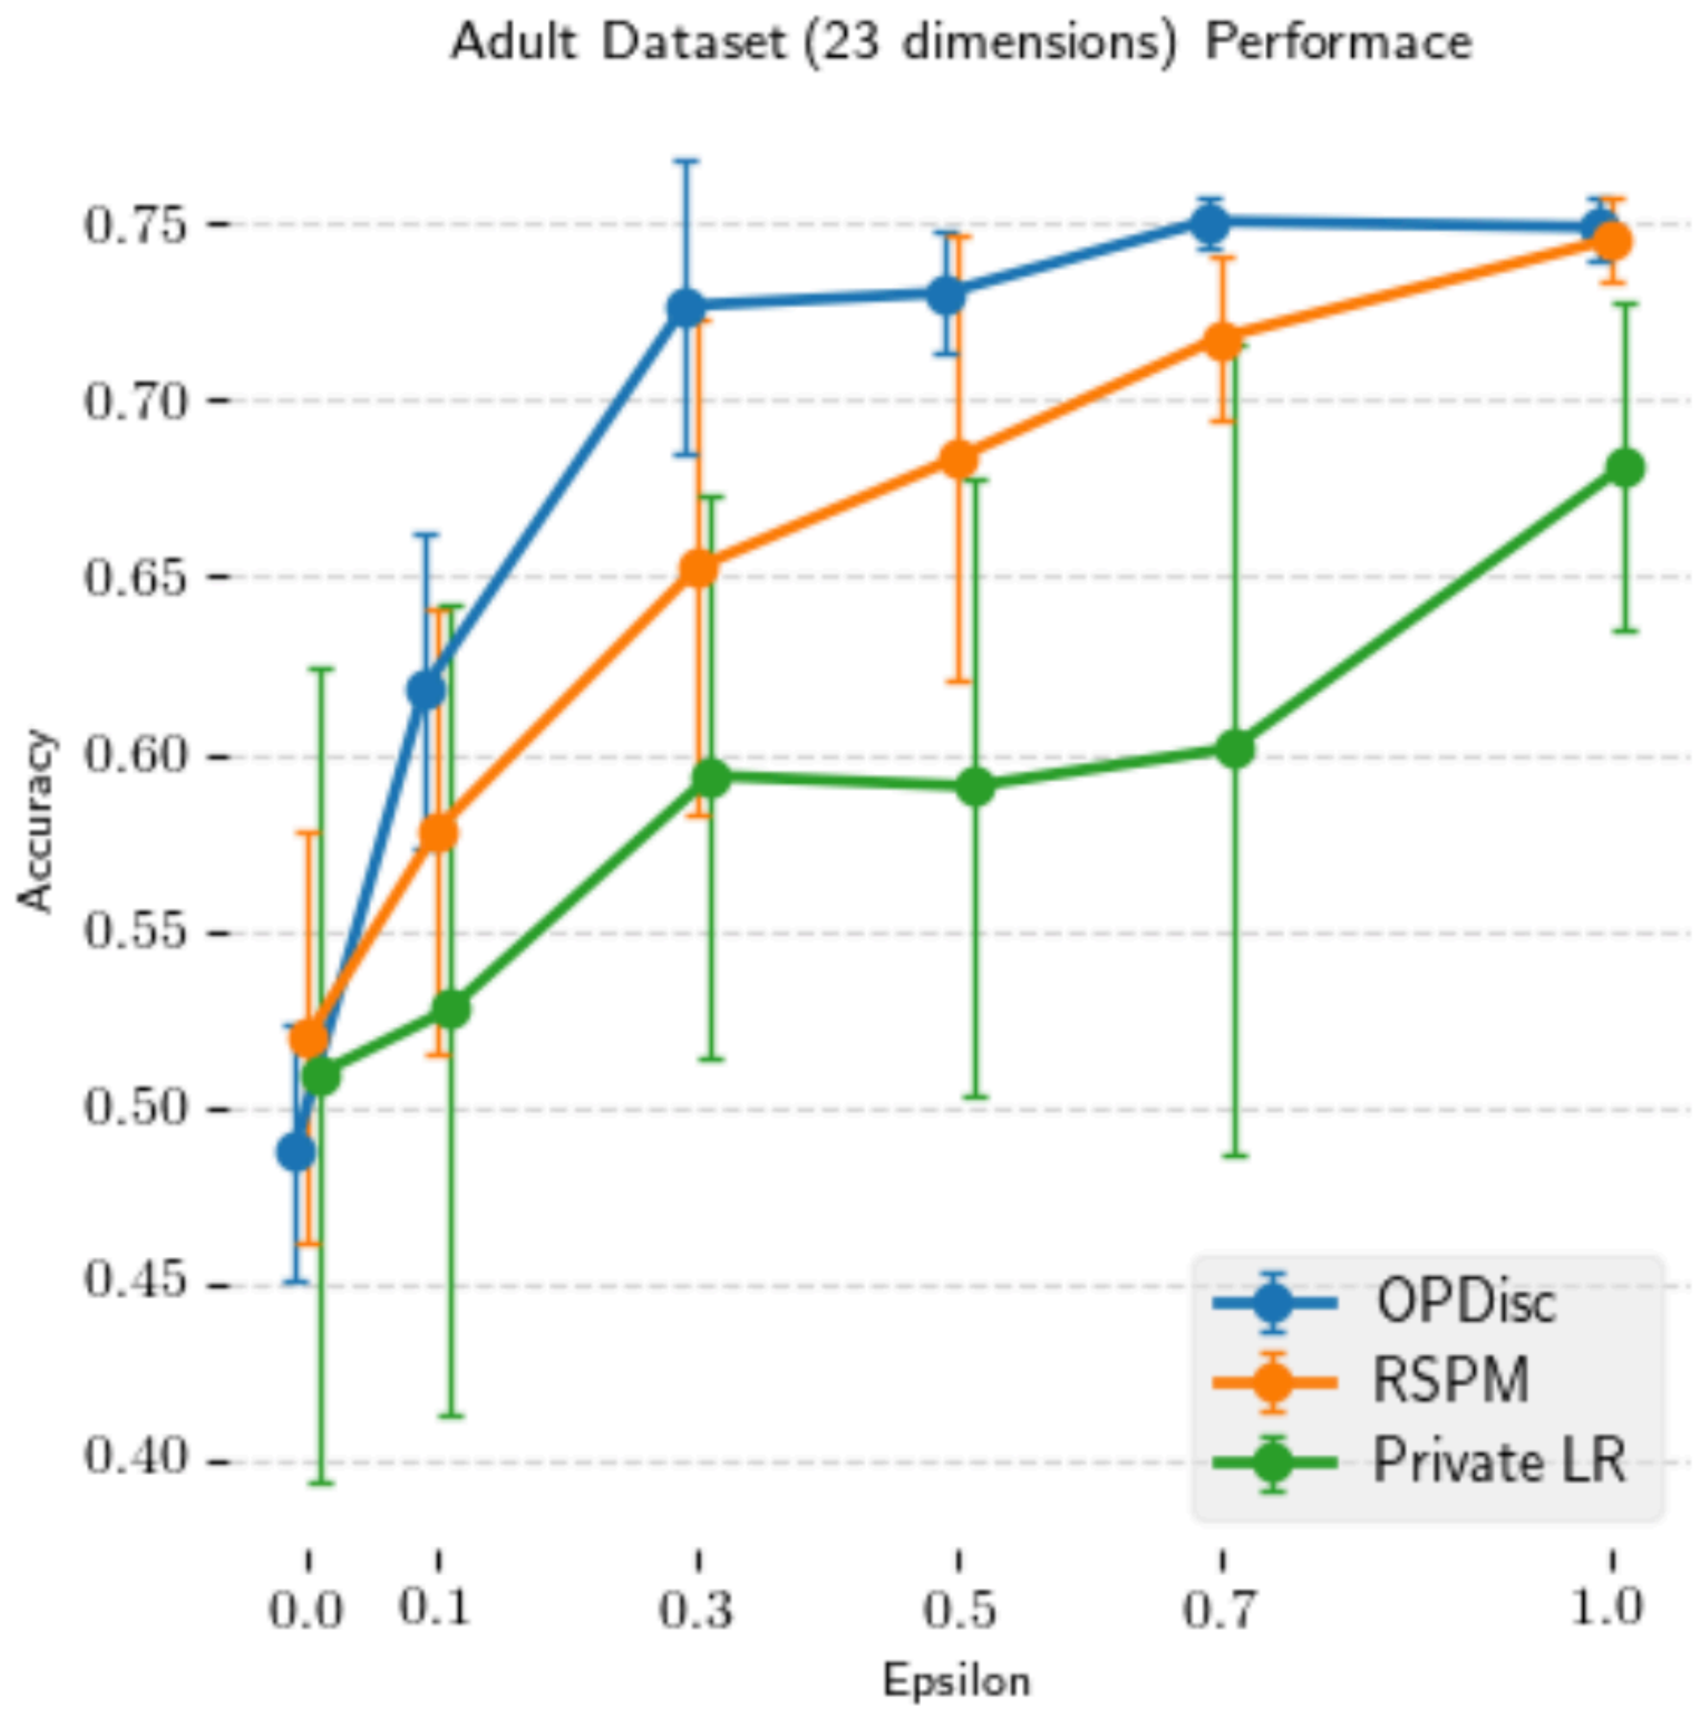
Adult Dataset (23 dimensions) Performace
Accuracy
0.75
0.70
0.65
0.60
0.55
0.50
0.45
0.40
0.0
0.1
0.3
0.5
0.7
1.0
Epsilon
OPDisc
RSPM
Private LR

Conclusion.

- We provide two new private algorithm,