Differentially Private Contextual Dynamic Pricing

Wei Tang (WashU)

Joint work with CJ Ho (WashU), Yang Liu (UCSC)

Differential Privacy (DP)

- **Neighboring Dataset:** $X', X \subset \mathbb{R}^d$ are neighbors if they differ in only one data of an individual.

- **Differential Privacy:** A randomized mechanism $\mathcal{M}: X \to O$ is ε-DP if for all neighboring inputs X', X, for all outputs $o \in O$ we have:

$$\mathbb{P}(\mathcal{M}(X) = o) \leq e^{\epsilon}\mathbb{P}(\mathcal{M}(X') = o)$$

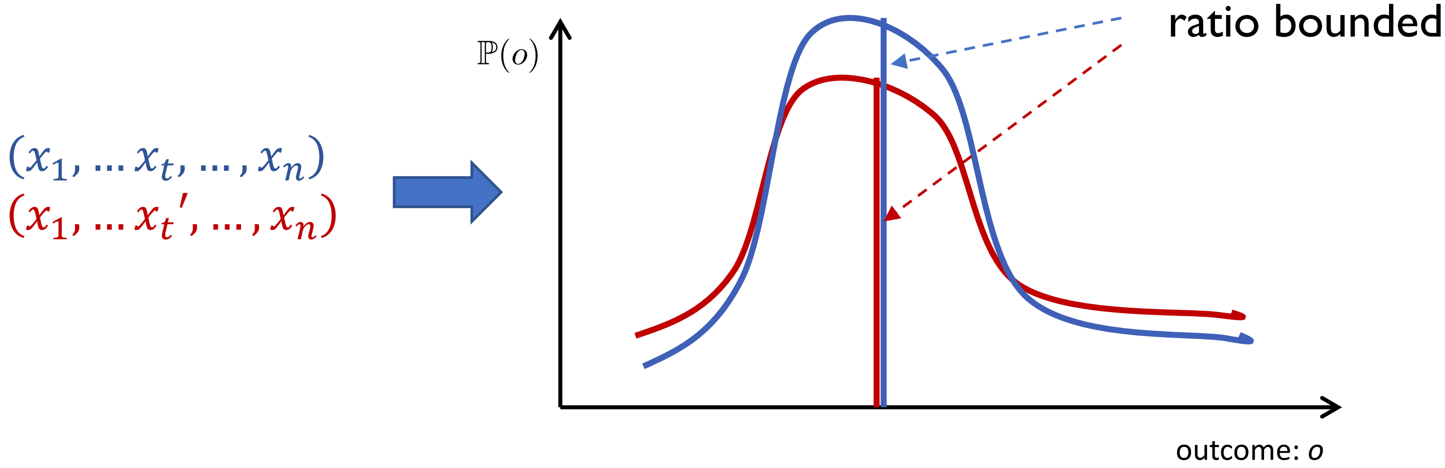

Differential Privacy (DP)

- **Neighboring Dataset:** $X', X \subset \mathbb{R}^d$ are neighbors if they differ in only one data of an individual.

- **Differential Privacy:** A randomized mechanism $\mathcal{M}: X \rightarrow O$ is ε-DP if for all neighboring inputs X', X, for all outputs $o \in O$ we have:

$$\mathbb{P}(\mathcal{M}(X) = o) \leq e^{\epsilon}\mathbb{P}(\mathcal{M}(X') = o)$$

  - ε smaller, strongly privacy guarantee
  - For small ε: $e^{\varepsilon} \approx 1 + \varepsilon \approx 1$

Bound the "maximum amount" that one person's data can change the output of a computation.

Some Key Properties of DP

- **Robustness to post-processing:** If $\mathcal{M}: \mathbb{R}^d \rightarrow O$ is ε-DP, then for any arbitrary randomized mapping $F: O \rightarrow O',$ the mechanism $F \circ \mathcal{M}$ is also ε-DP

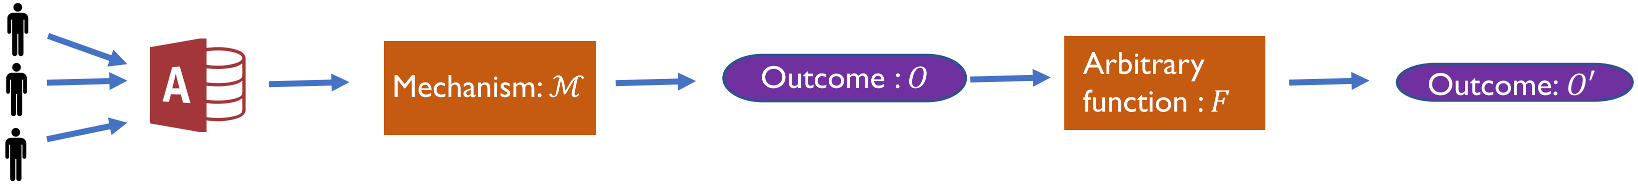

Some Key Properties of DP

- Robustness to post-processing: If $\mathcal{M}: \mathbb{R}^d \to O$ is ε-DP, then for any arbitrary randomized mapping $F: O \to O'$, the mechanism $F \circ \mathcal{M}$ is also ε-DP

- Composition: For $j \in [k]$, if $\mathcal{M}_j$ is ε_j - DP, then the mechanism $(\mathcal{M}_1, \ldots, \mathcal{M}_k)$ is $\sum_j \varepsilon_j$ - DP

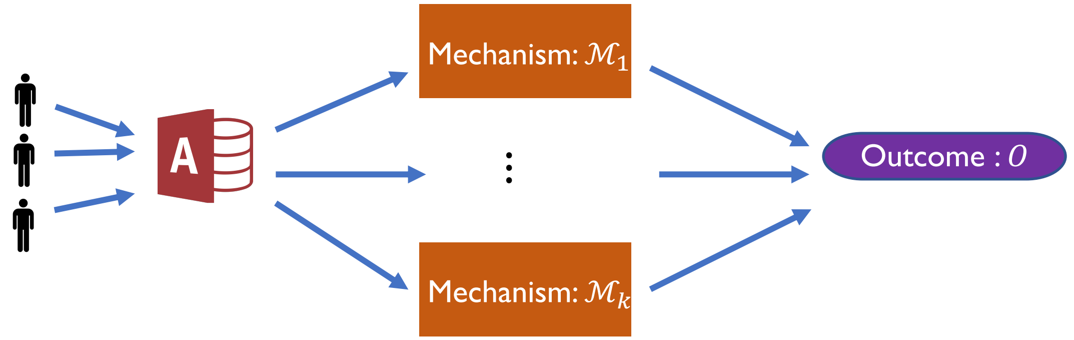

Contextual Dynamic Pricing

- In each timestep: *seller* has a *good* to sell to a *buyer* and needs to decide which price to put it in the market.
- At each time step t:
  - Seller receives a good $\boldsymbol{x}_t \in \mathbb{R}^d$
  - Buyer's value v_t: unknown to seller
  - Seller sets a price p_t and observes $y_t = \mathbb{1}_{\{p_t \leq v_t\}}$:
    - ➢ $p_t \leq v_t$, a sale is achieved and seller collects revenue $r_t = p_t$;
    - ➢ $p_t > v_t$, no sale is achieved and seller collects zero revenue: $r_t = 0$.
- Applications: online advertisements; real-estate,

Contextual Dynamic Pricing

- **Privacy Leakage**
  - Optimal Pricing policy is possible!
  - Buyers' past purchases are sensitive personal information.

- Goal: design a pricing policy which not only maximize her revenue but also protect the buyers' personal information

Private Pricing -- Objective

Privacy Guarantee

- Use differential privacy as privacy measure.
- A pricing policy $\mathcal{A}$
  - Feature vector sequence: $\boldsymbol{X} = \{\boldsymbol{x}_t\}_{t\geq 1}$;
  - Valuation sequence: $\boldsymbol{V} = \{v_t\}_{t\geq 1}$;
  - Response sequence: $\boldsymbol{Y} = \{y_t\}_{t\geq 1}$;
  - Price sequence: $\boldsymbol{P} = \{p_t\}_{t\geq 1}$

$$\Pr(\mathcal{A}(\boldsymbol{X}, \boldsymbol{Y} | \boldsymbol{V}) = \boldsymbol{P}) \leq e^{\varepsilon} \Pr(\mathcal{A}(\boldsymbol{X}, \boldsymbol{Y}' | \boldsymbol{V}') = \boldsymbol{P}) + \delta, \quad \forall \boldsymbol{P}$$

Private Pricing -- Objective

Utility Guarantee – minimize seller's Regret

$$\underbrace{\sum_{t=1}^{T} p_t^* \mathbb{I}_{\{p_t^* \le v_t\}}}_{\text{OPT}} - \underbrace{\sum_{t=1}^{T} p_t \mathbb{I}_{\{p_t \le v_t\}}}_{\text{Performance}}$$

- p_t^*: optimal price for good x_t -- knows the hidden v_t

Private Pricing -- Objective

Utility Guarantee – minimize seller's Regret

$$\text{Regret}_{\mathcal{A}}(T) = \sup_{X} \left(\underbrace{\sum_{t=1}^{T} p_t^* \mathbb{I}_{\{p_t^* \le v_t\}}}_{\text{OPT}} - \underbrace{\sum_{t=1}^{T} p_t \mathbb{I}_{\{p_t \le v_t\}}}_{\text{Performance}} \right)$$

For any adversarial arrival products

- p_t^*: optimal price for good x_t -- knows the hidden v_t

Sublinear regret: $\text{Regret}_{\mathcal{A}}(T) = o(T)$

Assumptions We Make

To solve the problem, we assume:

- Linear valuation : $v_t(\boldsymbol{x}_t) = \boldsymbol{\theta}^\top \boldsymbol{x}_t + z_t$
  - $\boldsymbol{\theta}$: unknown but fixed;
  - $z_t \sim F$: i.i.d drawn from F
  - By Postprocessing property, protecting $\{v_t\}$ reduce to protect $\{z_t\}$

- $F(v)$ and $1 - F(v)$ are log-concave in v.
  - A function f is log-concave $\rightarrow \log f$ is concave.
  - Including normal, uniform, and (truncated) Laplace, exponential, and logistic distributions.

Main Results

Main Result: reduction to online convex optimization with desired privacy Guarantee.

- $\text{Regret}_{\mathcal{A}}(T) = \sup_{\boldsymbol{X}}\left(\sum_{t=1}^{T} p_t^* \mathbb{I}_{\{p_t^* \leq v_t\}} - \sum_{t=1}^{T} p_t \mathbb{I}_{\{p_t \leq v_t\}}\right)$ protect $\{z_t\}$.

non-convex and no first order information

$l_t(\boldsymbol{\theta}) = -\mathbb{I}_{\{p_t \leq v_t\}} \log(1 - F(p_t - \langle \boldsymbol{x}_t, \boldsymbol{\theta} \rangle)) - \mathbb{I}_{\{p_t > v_t\}} \log(F(p_t - \langle \boldsymbol{x}_t, \boldsymbol{\theta} \rangle))$: **Convex!**

- $\text{Regret}_{\mathcal{A}}^{\boldsymbol{\theta}}(T) = \sup_{X} \sum_{t=1}^{T}\left(l_t(\widehat{\boldsymbol{\theta}}_t) - l_t(\boldsymbol{\theta})\right)$ protect $\{\widehat{\boldsymbol{\theta}}_t\}$.

Main Results

Main Result: reduction to online convex optimization with desired privacy Guarantee.

- $\text{Regret}^{\boldsymbol{\theta}}_{\mathcal{A}}(T) = \sup_{\boldsymbol{X}} \sum_{t=1}^{T} \left(l_t(\widehat{\boldsymbol{\theta}}_t) - l_t(\boldsymbol{\theta}) \right)$ protect $\{\widehat{\boldsymbol{\theta}}_t\}$.

Theorem: We can design an algorithm which achieves regret of $\tilde{O}(\sqrt{dT}/\varepsilon)$ with ensuring it is ε-differentially private.

d = feature dimensions, T = number of arrivals, $\tilde{O}$ suppress the logarithmic factors

- Note: the best-known bound of non-private policy's is $\tilde{O}(\sqrt{T})$
- Only worse to constant factor $\sqrt{d}/\varepsilon$

Technique

Main Result: reduction to online convex optimization with desired privacy Guarantee.

- $\text{Regret}_{\mathcal{A}}^{\boldsymbol{\theta}}(T) = \sup_{\boldsymbol{X}} \sum_{t=1}^{T} \left(l_t(\widehat{\boldsymbol{\theta}}_t) - l_t(\boldsymbol{\theta}) \right)$ protect $\{\widehat{\boldsymbol{\theta}}_t\}$.
- Online gradient descent doesn't work: $\widehat{\boldsymbol{\theta}}_{t+1} = \widehat{\boldsymbol{\theta}}_t - \eta_t \nabla l_t(\widehat{\boldsymbol{\theta}}_t)$
  - By post-processing property: reduce to ensure $\mathcal{A}$ is ε-DP w.r.t sequences of $\left(\nabla l_1(\widehat{\boldsymbol{\theta}}_1), \nabla l_2(\widehat{\boldsymbol{\theta}}_2), \ldots, \nabla l_T(\widehat{\boldsymbol{\theta}}_T) \right)$

Technique

- Prototypical algorithms for online convex optimization
    - Gradient Descent: $\widehat{\boldsymbol{\theta}}_{t+1} = \text{Project}_{\Theta}\left(\widehat{\boldsymbol{\theta}}_t - \eta \nabla l_t(\widehat{\boldsymbol{\theta}}_t)\right)$

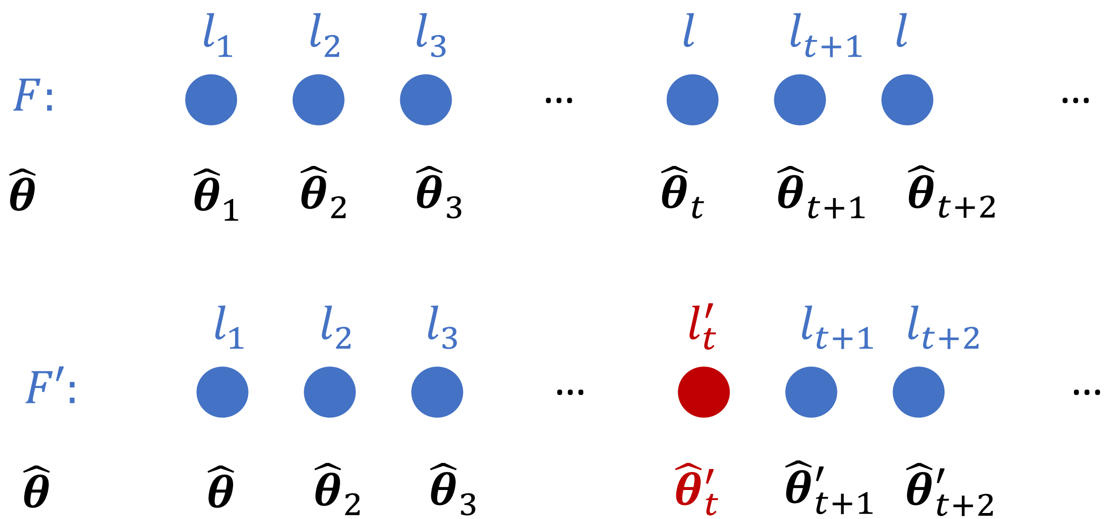

One single change in F will influence all subsequent updates on $\widehat{\boldsymbol{\theta}}$, which exaggerate the added noise to ensure privacy!

Technique

- Follow The Approximate Leader (FTAL)
  - Use all previous $\{\widehat{\boldsymbol{\theta}}_s\}_{s<t}$ to compute the $\widehat{\boldsymbol{\theta}}_t$
  - $\widehat{\boldsymbol{\theta}}_t = \mathrm{argmax}_{\widehat{\boldsymbol{\theta}}\in\Theta}\langle \sum_{s=1}^{t-1} \nabla l_s(\widehat{\boldsymbol{\theta}}_s),\, \widehat{\boldsymbol{\theta}} \rangle$

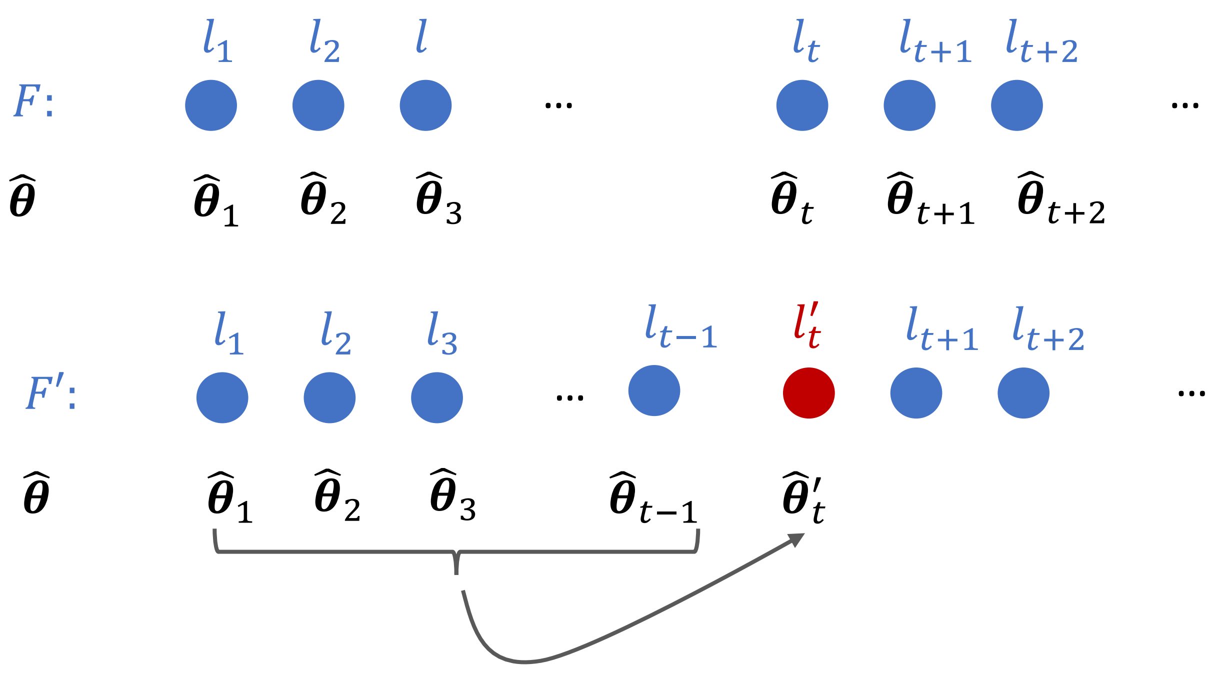

Technique

- Private FTAL: $\widehat{\boldsymbol{\theta}}_t = \text{argmax}_{\widehat{\boldsymbol{\theta}} \in \Theta} \langle \sum_{s=1}^{t-1} \nabla l_s(\widehat{\boldsymbol{\theta}}_s), \widehat{\boldsymbol{\theta}} \rangle$
  - $\sum_{s=1}^{t-1} \nabla l(\widehat{\boldsymbol{\theta}}_s)$ is DP
  - Tree-based Aggregation Protocol on high-dimensional space

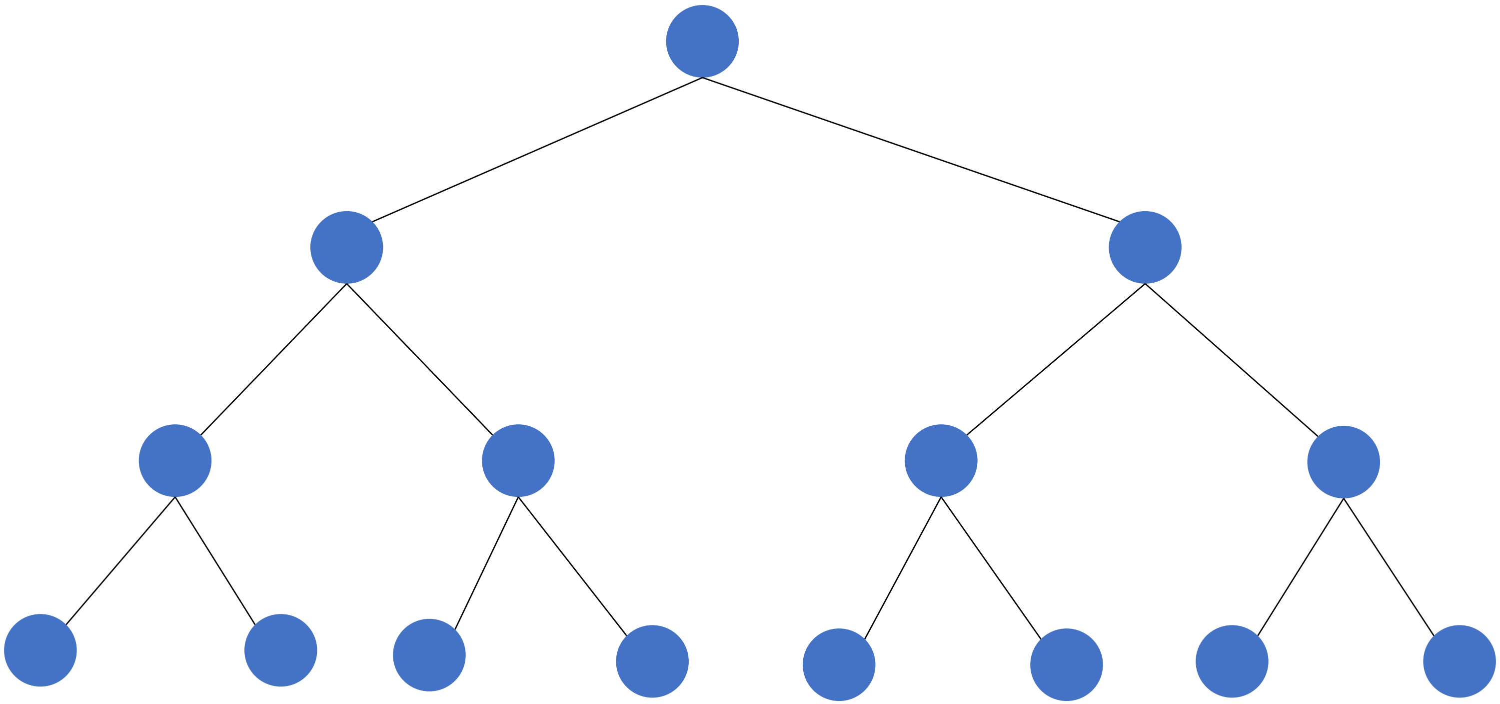

Data: $\nabla l_1(\widehat{\boldsymbol{\theta}}_1)$ $\nabla l_2(\widehat{\boldsymbol{\theta}}_2)$ $\nabla l_3(\widehat{\boldsymbol{\theta}}_3)$ $\nabla l(\widehat{\boldsymbol{\theta}}_4)$ $\nabla l_5(\widehat{\boldsymbol{\theta}}_5)$ $\nabla l_6(\widehat{\boldsymbol{\theta}}_6)$ $\nabla l_7(\widehat{\boldsymbol{\theta}}_7)$ ...

Thank you.