Private Selection from Private Candidates

Jingcheng Liu (University of California, Berkeley ⟶ Caltech)
Kunal Talwar (Google AI)

Private Selection

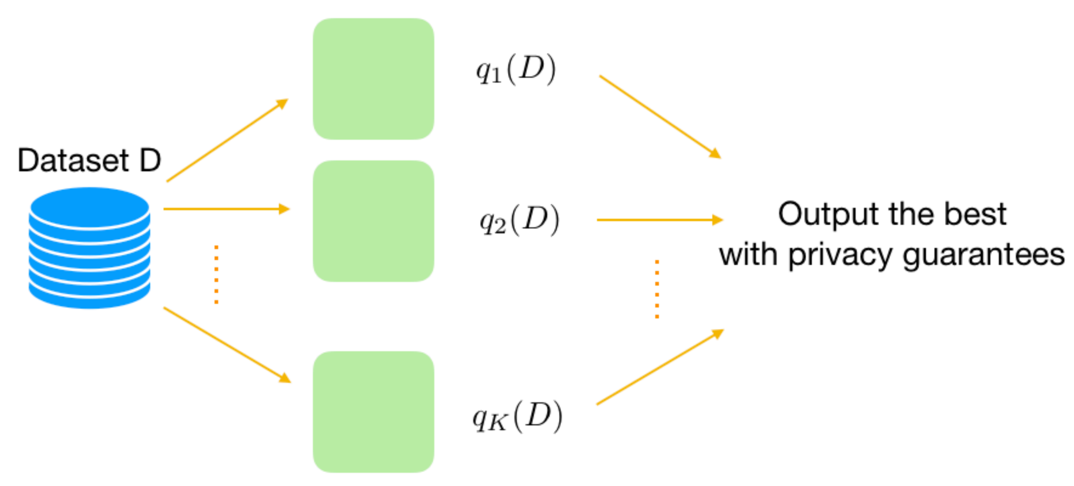

Life is all about choices!

Differential privacy (DP)

(by Dwork, McSherry, Nissim, Smith)

Whether you are in the data set or not, it makes little difference

Let $\mathcal{M}: \mathcal{D}^n \to \mathcal{R}$ be a randomized algorithm

We say that $\mathcal{M}$ satisfies $(\varepsilon, \delta) - \mathrm{DP}$ if $\forall D, D'$ s.t. $|D - D'| \le 1, \forall S$

$$\Pr[\mathcal{M}(D) \in S] \le \exp(\varepsilon) \cdot \Pr[\mathcal{M}(D') \in S] + \delta.$$

If $\delta = 0$, we say that $\mathcal{M}$ satisfies $\varepsilon - \mathrm{DP}$

Composition Theorems

$\Pr[\mathcal{M}(D) \in S] \le e^{\varepsilon} \cdot \Pr[\mathcal{M}(D') \in S]$

- Basic composition

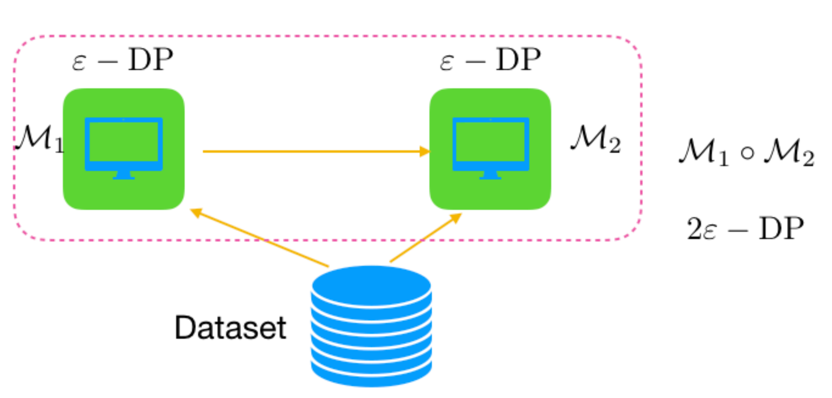

K-fold composition: $(K\varepsilon) - \mathrm{DP}$

- Advanced composition: $\left(\sqrt{2K \ln \frac{1}{\delta}} \cdot \varepsilon + 2k\varepsilon^2, \delta\right) - \mathrm{DP}$

Private Selection for Lipschitz Functions

Exponential mechanism [McSherry, Talwar]

If $\forall i, q_i(D)$ is $1 - \text{Lipschitz}$: $|q_i(D) - q_i(D')| \le 1$

$$\max\{q_i(D) + \mathrm{Lap}(1/\varepsilon)\} \text{ is } \varepsilon - \mathrm{DP}$$

Privacy

Moreover, if the index j is the maximizer, then

$$\Pr\left[q_j(D) \le \max_i \{q_i(D)\} - \frac{1}{\varepsilon} \ln \frac{K}{\delta}\right] \le \delta$$

Utility

An example:

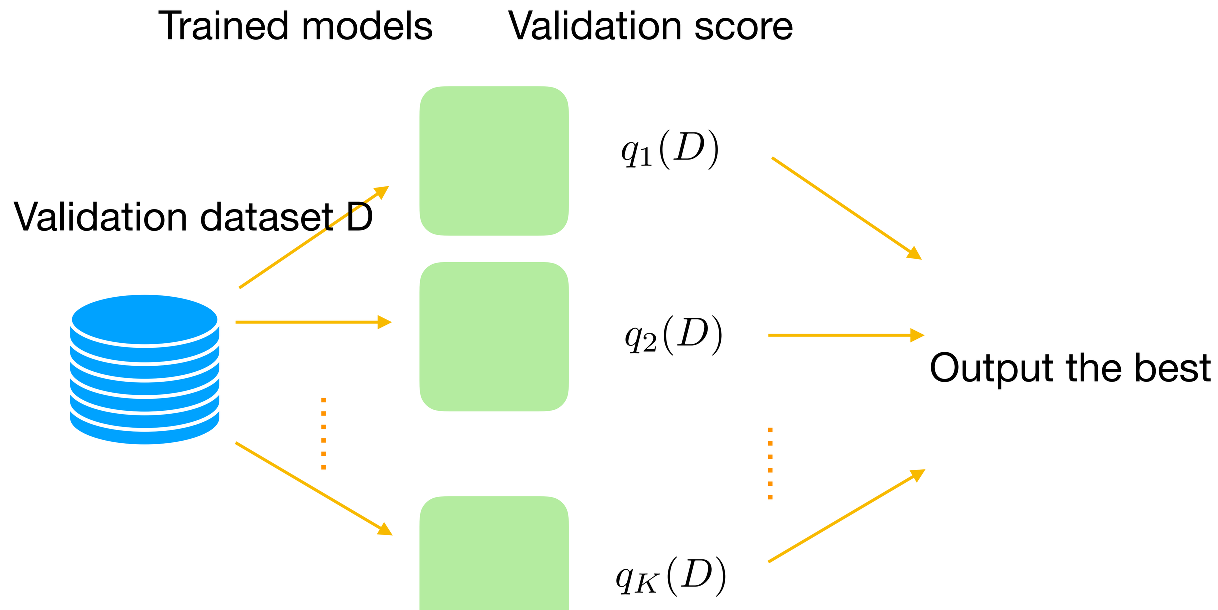

$q_i(D) :=$ (percentage of correctness of the i-th model)

$\forall i, q_i(D)$ is $\frac{1}{n} - \text{Lipschitz}$: $|q_i(D) - q_i(D')| \le 1/n$

Private Candidates

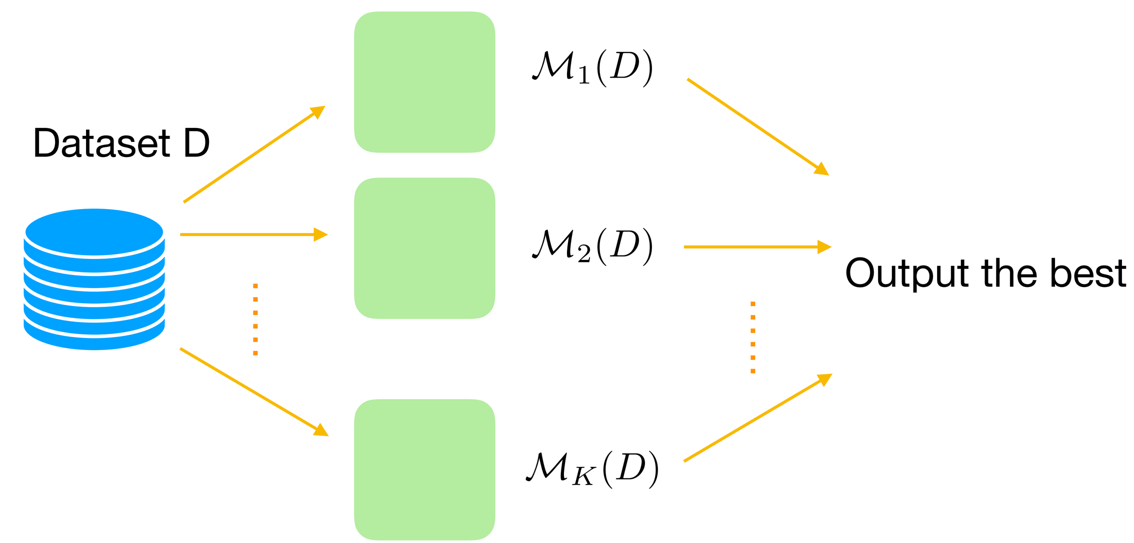

$(x_i, q_i) \leftarrow \mathcal{M}_i(D)$, which are $\varepsilon - \mathrm{DP}$ randomized algorithms

output score

Some *real* examples

- Algorithm selection
- Model selection
- Neural network architecture
- Hyperparameters selection
- ....

Private Selection: Naive Attempt #1

Since $\forall i, \mathcal{M}_i(D)$ is $\varepsilon - \mathrm{DP}$, what if we choose $\max_i \mathcal{M}_i$?

Basic composition: $(K\varepsilon) - \mathrm{DP}$

Tight example for attempt #1

$\Pr[\mathcal{M}(D) \in S] \le e^{\varepsilon} \cdot \Pr[\mathcal{M}(D') \in S]$

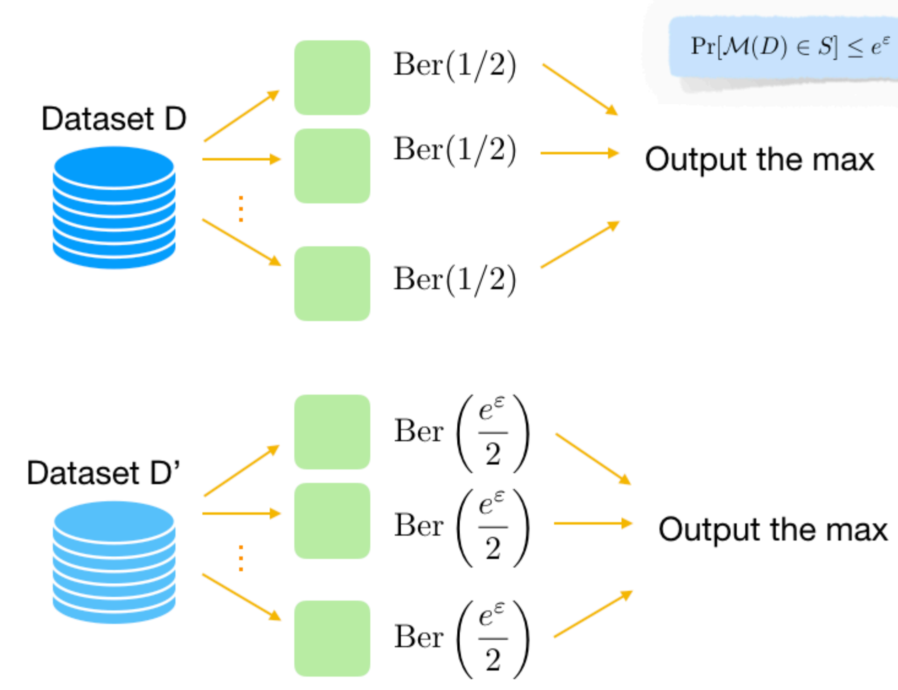

$\Pr[\text{output}(D) = 0] = \frac{1}{2^K}$, $\Pr[\text{output}(D') = 0] = \left(1 - \frac{e^{\varepsilon}}{2}\right)^K \approx \frac{e^{-K\varepsilon}}{2^K}$

Private Selection: Naive Attempt #2

Since $\forall i, \mathcal{M}_i(D)$ is $\varepsilon - \mathrm{DP}$, what if we choose $i \sim [K]$ u.a.r?

Indeed we do get $\varepsilon - \mathrm{DP}$, but the probability of getting the best can be $\frac{1}{K}$

Repetition with random stopping

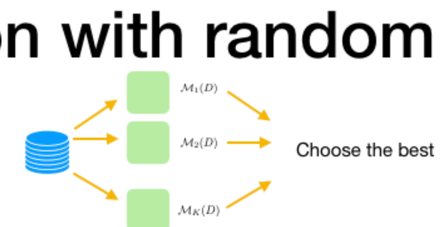

Algorithm 1:

Repeat for T rounds:

$Q(D)$:
1. Choose $i \sim [K]$ u.a.r
2. Get $(x_i, q_i) \leftarrow \mathcal{M}_i(D)$

$\varepsilon - \mathrm{DP}$

Output the best of seen-so-far

- For every fixed T, Algorithm 1 is $(T\varepsilon) - \mathrm{DP}$
- Our result: if $T \sim \mathrm{Geom}(\cdot)$, then Algorithm 1 is $(3\varepsilon) - \mathrm{DP}$

Remarks

- Generalizes 1-Lipschitz queries
- For a suitable choice of Geom(.), can match the utility of the Exp. Mech.

Prior work

- Assume Lipschitzness: Exponential mechanism [McSherry, Talwar' 07]
- Assume "local" Lipschitzness [Raskhodnikova, Smith' 16]
- Assume a known "good" target [Gupta, Ligett, McSherry, Roth, Talwar' 10]

We are able to improve upon privacy, utility, and computational (sampling) efficiency

Utility for random stopping

Fix any constants $\alpha \in (0,1), \eta \in (0, 1/5)$, and choose $\gamma = \frac{4\alpha\eta}{5 \ln 1/\eta}$ in Algorithm 1. If there is a threshold τ (unknown to the algorithm), and an event $\mathcal{G}$ (on the output of Q) such that

$$\Pr_{\tilde{q} \sim Q(D)} [\tilde{q} \ge \tau] \ge \alpha,$$

$$\Pr_{\tilde{q} \sim Q(D)} [\tilde{q} \ge \tau \wedge \bar{\mathcal{G}}] \le \frac{\alpha\eta^2}{\ln^2 \frac{1}{\eta}}.$$

Let $A_{\mathrm{out}}(D)$ be the output of Algorithm 1 on D. Then, we have $\Pr[A_{\mathrm{out}}(D) \in \mathcal{G}] \ge 1 - 5\eta$.

Adaptive Private Selection

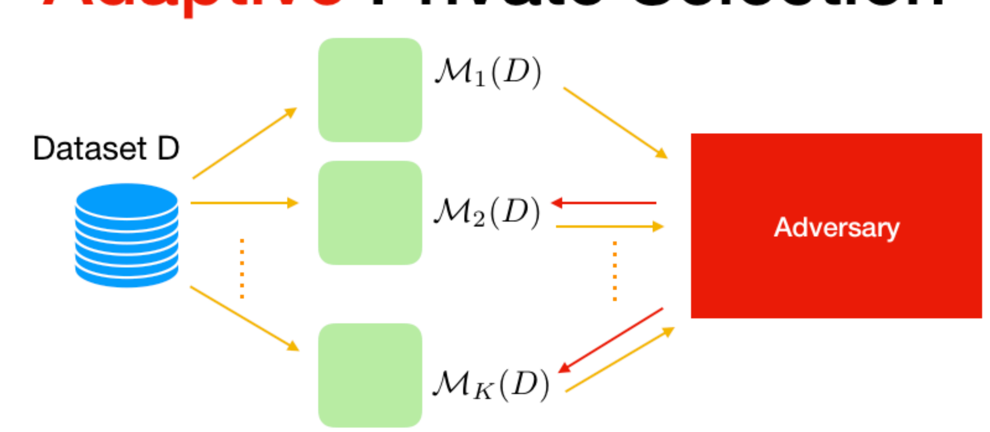

Actually, life is about adaptive choices….

Indeed, there is an adaptive version of the Exp. Mech.:
- the sparse vector algorithm
- due to Dwork, Naor, Reingold, Rothblum, and Vadhan '09

Motivating Applications

- Adaptive tuning in machine learning
- Adaptive data analysis: garden of forking paths. [Dwork, Feldman, Hardt, Pitassi, Reingold, and Roth' 15]
- …

Sparse vector algorithm only works for Lipschitz queries. Can we go beyond Lipschitz queries?

Adaptive Private Selection

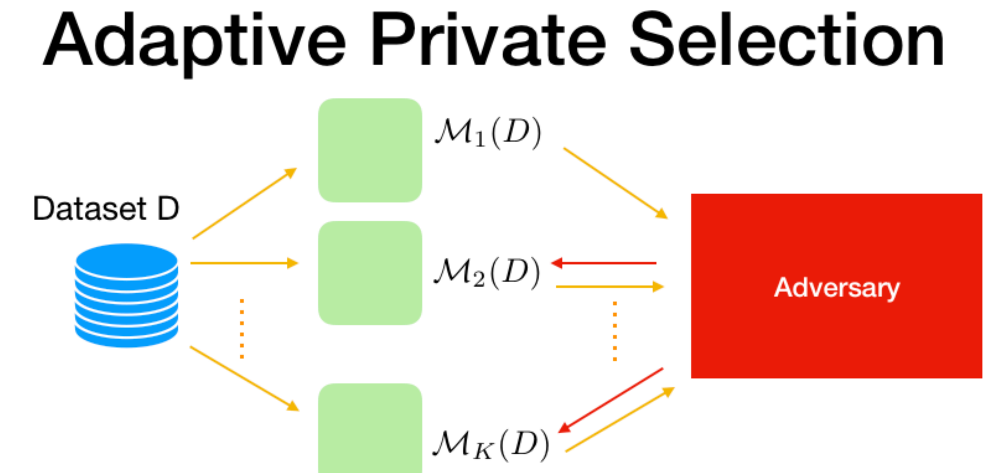

- Allow adaptive queries: $\mathcal{M}_i(D)$ are $\varepsilon - \mathrm{DP}$ randomized algorithms
- But only threshold queries

Median score

$$\mathrm{Median}(\mathcal{M}_i(D)) := \sup \left\{\tau : \Pr_{(\tilde{x}, \tilde{q}) \sim \mathcal{M}_i(D)} [\tilde{q} \ge \tau] \ge \frac{1}{2}\right\}.$$

Given a threshold τ
Goal: output the first i such that $\mathrm{Median}(\mathcal{M}_i(D)) \ge \tau$.

Our Result for Adaptive Private Selection:

There is an (ε, δ)-DP algorithm such that w.h.p.
- If a query is above the threshold, then Alg. reports "AboveThreshold"
- If Alg. reports i-th query is "AboveThreshold", then the i-th query is not "too much below the threshold" in terms of the "percentile-score"

Effectively, there is a reasonable probability of exceeding the threshold

Our approach:
- Introduce "percentile-score", which generalizes the median score
- Estimating & testing the "percentile-score" differentially privately
- Bound a variant of Earth mover's distance between sums of Bernoullis.

Open Problems & future work:

- Other stopping time distribution?
- DP preserving mechanism for more sophisticated (hyperparameter) optimization algorithm