CSC/DSCC 574 Continuous Algorithms for Optimization and Sampling Lecture 7

Differential Privacy

Lecturer: Jiaming Liang *February 22, 2024*

1 Differential privacy

Definition 1 ((ε, δ)-DP)**.** *A randomized mechanism $\mathcal{M}$ is (ε, δ)-differentially private if for any neighboring databases $\mathcal{D}, \mathcal{D}'$ and any subset $S \subseteq O$ (output space), one has*

$$\mathbb{P}[\mathcal{M}(\mathcal{D}) \in S] \leq e^{\varepsilon} \mathbb{P}\left[\mathcal{M}\left(\mathcal{D}'\right) \in S\right] + \delta.$$

We say $\mathcal{D}$ and $\mathcal{D}'$ are neighboring databases if they agree on all the user inputs except for a single user's input.

For $\delta = 0$, the ε-DP condition can be written as

$$\frac{1}{e^{\varepsilon}} \mathbb{P}\left[\mathcal{M}\left(\mathcal{D}'\right) \in S\right] \leq \mathbb{P}[\mathcal{M}(\mathcal{D}) \in S] \leq e^{\varepsilon} \mathbb{P}\left[\mathcal{M}\left(\mathcal{D}'\right) \in S\right].$$

1.1 Randomized response

Consider a survey at U of R where we want to know the percentage of students who cheated on an exam. Here is an idea that protects students' privacy while we still get useful information. First, the surveyee flips a fair coin. If it is tail, then answer the question truthfully. If it is head, flip again. Based on outcome of second coin flip, answer "YES" if it is head and "NO" if it is tail.

The appeal to tell the truth here is that if you answer "YES", you have plausible deniability for saying why this is true or not. The above process can be put into the DP setting. Input databases $\mathcal{D}_Y = \{Y\}$ and $\mathcal{D}_N = \{N\}$ and output subsets $S = \{Y\}$ or $S = \{N\}$. It is easy to see that

$$\mathbb{P}(Y|Y) = \frac{3}{4}, \quad , \mathbb{P}(N|Y) = \frac{1}{4}, \quad , \mathbb{P}(N|N) = \frac{3}{4}, \quad , \mathbb{P}(Y|N) = \frac{1}{4},$$

i.e.,

$$\mathbb{P}[\mathcal{M}(\mathcal{D}_Y) \in \{Y\}] = \frac{3}{4}, \quad \mathbb{P}[\mathcal{M}(\mathcal{D}_Y) \in \{N\}] = \frac{1}{4}, \quad \mathbb{P}[\mathcal{M}(\mathcal{D}_N) \in \{N\}] = \frac{3}{4}, \quad \mathbb{P}[\mathcal{M}(\mathcal{D}_N) \in \{Y\}] = \frac{1}{4}.$$

Thus, for any neighboring $\mathcal{D}$ and $\mathcal{D}'$, we have

$$\frac{1}{3} \leq \frac{\mathbb{P}[\mathcal{M}(\mathcal{D}) \in \{Y\}]}{\mathbb{P}[\mathcal{M}(\mathcal{D}') \in \{Y\}]} \leq 3, \quad \frac{1}{3} \leq \frac{\mathbb{P}[\mathcal{M}(\mathcal{D}) \in \{N\}]}{\mathbb{P}[\mathcal{M}(\mathcal{D}') \in \{N\}]} \leq 3.$$

Therefore for any S and $\mathcal{D}, \mathcal{D}'$, we have

$$\frac{1}{3} \leq \frac{\mathbb{P}[\mathcal{M}(\mathcal{D}) \in S]}{\mathbb{P}[\mathcal{M}(\mathcal{D}') \in S]} \leq 3.$$

The randomized response mechanism is $\ln 3$-DP

If surveyee follows the randomized response mechanism, then the expected faction of "YES" is $\frac{3}{4}p + \frac{1}{4}(1-p)$ where p is the true fraction. The empirical fraction $\hat{p}$ of "YES" converges to $p/2 + 1/2$ as $n \to \infty$, and we will get better estimate of p.

1.2 Laplace mechanism

Consider $\mathcal{M}(D) = \mathcal{A}(D) + X$ where $X \sim \text{Laplace}(0, b)$, then $\mathcal{M}(D) \sim \text{Laplace}(\mathcal{A}(D), b)$, i.e.,

$$p(\mathcal{M}(\mathcal{D}) = s) = \frac{1}{2b} \exp\left(-\frac{|s - \mathcal{A}(\mathcal{D})|}{b}\right).$$

Hence, we have

$$\begin{aligned}\frac{\mathbb{P}[\mathcal{M}(\mathcal{D}) \in S]}{\mathbb{P}[\mathcal{M}(\mathcal{D}') \in S]} &= \frac{\int_{s \in S} \frac{1}{2b} \exp\left(-\frac{|s-\mathcal{A}(\mathcal{D})|}{b}\right) ds}{\int_{s \in S} \frac{1}{2b} \exp\left(-\frac{|s-\mathcal{A}(\mathcal{D}')|}{b}\right) ds} \leq \max_{s \in S} \frac{\exp\left(-\frac{|s-\mathcal{A}(\mathcal{D})|}{b}\right)}{\exp\left(-\frac{|s-\mathcal{A}(\mathcal{D}')|}{b}\right)} \\ &= \max_{s \in S} \exp\left(\frac{|s - \mathcal{A}(\mathcal{D}')| - |s - \mathcal{A}(\mathcal{D})|}{b}\right) \leq \exp\left(\frac{|\mathcal{A}(\mathcal{D}') - \mathcal{A}(\mathcal{D})|}{b}\right).\end{aligned}$$

Suppose $\mathcal{D}$ is a collection of human heights and $\mathcal{A}(\mathcal{D})$ is the average, then we have a bound on $\Delta = |\mathcal{A}(\mathcal{D}') - \mathcal{A}(\mathcal{D})|$. In general, we assume Δ exists and choose $b = \Delta/\varepsilon$, then the Laplace mechanism is ε-DP, or $(\varepsilon, 0)$-DP.

1.3 Exponential mechnism

The exponential mechanism (EM) was designed for situations in which we wish to choose the "best" response but adding noise directly to the computed quantity can completely destroy its value, such as setting a price in an auction, where the goal is to maximize revenue, and adding a small amount of positive noise to the optimal price (in order to protect the privacy of a bid) could dramatically reduce the resulting revenue.

Another example is training a neural network given a database of training data. The neural network returned is defined by a series of weights. If we were to apply the Laplace Mechanism to this function, Laplace noise would be added to the weights before returning the network. However, even small fluctuations in weights in a neural network may severely impact the performance of that network. Therefore, the returned network (with added noise) will likely behave very differently than the initial network found (before adding noise) that minimized error, and thus would have a unpredictably higher error than the minimal error network we desired.

Given input database $\mathcal{D}$ and output $s \in O$, consider a utility function $u(\mathcal{D}, s)$. This utility function is the revenue in the auction and is the negative loss function in neural network training. Then the exponential mechanism is $\mathcal{M}_u(\mathcal{D}) \sim \pi(\cdot) \propto \exp\left(ku(\mathcal{D}, \cdot)\right)$, i.e.,

$$p(\mathcal{M}_u(\mathcal{D}) = s) = \frac{\exp\left(ku(\mathcal{D}, s)\right)}{\int_{s' \in O} \exp\left(ku(\mathcal{D}, s')\right) ds'}.$$

Thus, we have

$$\mathbb{P}[\mathcal{M}_u(\mathcal{D}) \in S] = \int_{s \in S} \frac{\exp\left(ku(\mathcal{D}, s)\right)}{\int_{s' \in O} \exp\left(ku(\mathcal{D}, s')\right) ds'} ds.$$

It follows that

$$\begin{aligned}
\frac{\mathbb{P}[\mathcal{M}_u(\mathcal{D}) \in S]}{\mathbb{P}[\mathcal{M}_u(\mathcal{D}') \in S]} &= \frac{\int_{s \in S} \frac{\exp(ku(\mathcal{D}, s))}{\int_{s' \in O} \exp(ku(\mathcal{D}, s')) ds'} ds}{\int_{s \in S} \frac{\exp(ku(\mathcal{D}', s))}{\int_{s' \in O} \exp(ku(\mathcal{D}', s')) ds'} ds} \\
&= \frac{\int_{s \in S} \exp\left(ku(\mathcal{D}, s)\right) ds}{\int_{s \in S} \exp\left(ku(\mathcal{D}', s)\right) ds} \cdot \frac{\int_{s' \in O} \exp\left(ku(\mathcal{D}', s')\right) ds'}{\int_{s' \in O} \exp\left(ku(\mathcal{D}, s')\right) ds'} \\
&\leq \max_{s \in S} \exp\left(k[u(\mathcal{D}, s) - u(\mathcal{D}', s)]\right) \cdot \max_{s' \in O} \exp\left(k[u(\mathcal{D}', s') - u(\mathcal{D}, s')]\right) \\
&\leq \exp(2k\Delta_u)
\end{aligned}$$

where

$$\Delta_u := \max_{s \in O} \max_{D, D'} |u(\mathcal{D}, s) - u(\mathcal{D}', s)|.$$

If we choose $k = \varepsilon/(2\Delta_u)$, then the exponential methanism is $(\varepsilon, 0)$-DP.

2 Private convex optimization

Recall stochastic optimization is

$$\min_{x \in Q} \{f(x) = \mathbb{E}_\xi[F(x; \xi)]\},$$

and its SAA is

$$f(x; \mathcal{D}) = \frac{1}{n} \sum_{i=1}^{n} F(x; \xi_i),$$

where $\mathcal{D} = \{\xi_1, \ldots, \xi_n\}$ is a database.

Here $f(x; D)$ can be understood as the negative utility function $-u(D; s)$, where $x = s$ is the output of a certain mechanism We assume $F(\cdot; \xi)$ is convex and M-Lipschitz continuous, Q has a diameter $D > 0$.

We want to output a solution x^{priv} using a differentially private mechanism $\mathcal{M}$ such that we minimize the excess empirical risk

$$\mathbb{E}_{\mathcal{M}}\left[f\left(x^{\text{priv}};\mathcal{D}\right)\right] - f\left(x_*;\mathcal{D}\right),$$

where $x_* \in Q$ is the minimizer of $f(x;\mathcal{D})$.

In the literature, it is shown that EM achieves the optimal excess empirical risk $\mathcal{O}\left(\frac{MDd}{n\varepsilon}\right)$ under $(\varepsilon, 0)$-DP. On the other hand, it has also been shown that noisy gradient descent achives and excess empirical risk of

$$\mathcal{O}\left(\frac{MD\sqrt{d\log\frac{1}{\delta}}}{n\varepsilon}\right)$$

under (ε, δ)-DP, which is also shown to be optimal.

Note that the second bound only loses a lit in the privacy (δ) but reduces the dependence of d in the xcess empirical risk from d to $\sqrt{d}$. It is natural to ask the question whether we can obtain the optimal empirical risk under (ε, δ)-DP using EM. The answer is confirmative, but we need to introduce a modified version of EM, that is the regularized exponential mechanism,

$$x^{\text{priv}} \sim \exp\left(-kf(x;\mathcal{D}) + \frac{\mu}{2}\|x\|_2^2\right).$$

With a suitable choice of μ and k, we recover the optimal excess risk under (ε, δ)-DP.

EM is the task of sampling and the regularized EM is the core of proximal sampling. Before we are able to establish the aforementioned optimal excess risk under DP constraint, we first need to develop the algorithmic toolbox for sampling.