CSIS0351/CSIS8601: Randomized Algorithms
Lecture 9: Differential Privacy, Laplace Distribution
Instructor: Hubert Chan
Date: 14 Nov 2011

These lecture notes are supplementary materials for the lectures. They are by no means substitutes for attending lectures or replacement for your own notes!

1 Differential Privacy

1.1 Motivation

A *statistical database* is a database used for statistical analysis. For example, a database containing information about graduates of a certain university can answer questions like: what is the average salary of the graduates? Statistical database are widely used due to the enormous social value they provide: the previously mentioned database benefits the society in helping students to choose whether to go to that university, or how funding should be distributed among universities. However, the statistics released might cause leakage of sensitive information. Therefore, a big challenge is to maintain individual privacy, while providing useful aggregate statistical information about a certain group.

In 1977 the statistician Tore Dalenius gave a privacy goal for statistical databases: anything that can be learned about a member in the statistical database, should also can be learned without access to the database. However, as the following example illustrates, as long as the statistical database is useful, the goal is not achievable if the adversary has auxiliary information (information not obtained from the statistical database). Suppose we know that some student's X salary is \$10,000 more than the average, we can know his/her salary by querying about the average salary. Note that in this case, even if X does not join the database, we can still approximately know the average salary and hence know his/her sensitive information.

Notation. Let $\mathcal{U}$ be the set of possible user data points. A database of n users contains the data points for each user and can be viewed as a vector in $\mathcal{U}^n$. We use $\mathcal{D} := \mathcal{U}^n$ to denote the collection of all possible databases with n users. We consider whether releasing the output of some function $f : \mathcal{D} \to \mathcal{O}$ will compromise an individual's privacy.

Example. Suppose $\mathcal{U} = [0, 10^5]$ is the range of possible monthly salaries for a graduate. Given a database $X \in \mathcal{U}^n$, suppose the function of interest is $\mathsf{sum}(X) = \sum_{i=1}^{n} X_i$. At first sight, it might seem that releasing the sum does not violate an individual's privacy, because the sum does not directly reveal any individual's salary. However, in reality, users' data can be in many databases. Suppose users 1 to $n-1$ also participate in another database which also releases the sum. Then, combining the results of the two sums, the salary of user n can be accurately calculated!

Differential privacy defines privacy in a different sense: to minimize the increased risk of the sensitive information leakage due to one's joining in the statistical database. A private mechanism that answers queries to statistical databases can achieve this goal by introducing randomness such that when two databases differ by only one single user, the output produced have similar distributions.

Note that a differentially private mechanism encourages individuals to participate in statistical databases, and thus enhancing the social benefit provided by them.

1.2 Formal Definition

Let $\mathcal{U}$ be the set of possible data points and $\mathcal{D} := \mathcal{U}^n$ be the collection of databases with n users. Two databases $X^{(1)} \in \mathcal{D}$ and $X^{(2)} \in \mathcal{D}$ are called *neighboring* (denoted by $X^{(1)} \sim X^{(2)}$), if they differ by at most one coordinate.

Definition 1.1 (ϵ-differential privacy) *A randomized mechanism (function) $\mathcal{M} : \mathcal{D} \to \mathcal{O}$ preserves ϵ-differential privacy, if for any two neighboring databases $X^{(1)} \sim X^{(2)}$, and any possible set of output $\mathcal{S} \subseteq \mathcal{O}$, the following hold:*

$$\Pr[\mathcal{M}(X^{(1)}) \in \mathcal{S}] \leq \exp(\epsilon) \cdot \Pr[\mathcal{M}(X^{(2)}) \in \mathcal{S}],$$

where the randomness comes from the coin flips of $\mathcal{M}$.

Remark 1.2 We observe the following.

1. Since $\mathcal{M}$ is a randomized mechanism, $\mathcal{M}(X)$ is a distribution of outputs in $\mathcal{O}$.
2. Interchanging the roles of $X^{(1)}$ and $X^{(2)}$, we also have:
   $\Pr[\mathcal{M}(X^{(2)}) \in \mathcal{S}] \leq \exp(\epsilon) \cdot \Pr[\mathcal{M}(X^{(1)}) \in \mathcal{S}]$.
3. If $\mathcal{O}$ is a countable set, then we can also have the inequality for each $x \in \mathcal{O}$,
   $\Pr[\mathcal{M}(X^{(1)}) = x] \leq \exp(\epsilon) \cdot \Pr[\mathcal{M}(X^{(2)}) = x]$.
4. The inequality means that the distributions $\mathcal{M}(X^{(1)})$ and $\mathcal{M}(X^{(2)})$ are close, and hence by observing the output, it is not possible to tell for certain whether the output comes from database $X^{(1)}$ or $X^{(2)}$.

1.3 Properties of differentially private mechanisms

Theorem 1.3 *Let $\mathcal{M}_1 : \mathcal{D} \to \mathcal{O}_1$ be an ϵ_1-differentially private mechanism, and let $\mathcal{M}_2 : \mathcal{D} \to \mathcal{O}_2$ be an ϵ_2-differentially private mechanism. Also suppose $\mathcal{M}_1$ and $\mathcal{M}_2$ are independent. Then, the mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{O}_1 \times \mathcal{O}_2$ such that for any $X \in \mathcal{D}$, $\mathcal{M}(X) := (\mathcal{M}_1(X), \mathcal{M}_2(X))$, is $(\epsilon_1 + \epsilon_2)$-differentially private.*

Proof: Let $X_1 \in \mathcal{D}$ and $X_2 \in \mathcal{D}$ be two neighboring databases, and let $\mathcal{S} \subset \mathcal{O}_1 \times \mathcal{O}_2$ be a measurable set. For any $z \in \mathcal{O}_1$, define $\mathcal{S}_z := \{y \in O_2 : (z, y) \in \mathcal{S}\}$. Then, we have

$$\begin{aligned}
\Pr[\mathcal{M}_1(X_1) = z \wedge \mathcal{M}_2(X_1) \in \mathcal{S}_z] &= \Pr[\mathcal{M}_1(X_1) = z] \cdot \Pr[\mathcal{M}_2 \in \mathcal{S}_z] \\
&\leq \exp(\epsilon_1) \cdot \Pr[\mathcal{M}_1(X_2) = z] \cdot \exp(\epsilon_2) \cdot \Pr[\mathcal{M}_2(X_2) \in \mathcal{S}_z] \\
&= \exp(\epsilon_1 + \epsilon_2) \cdot \Pr[\mathcal{M}_1(X_2) = z \wedge \mathcal{M}_2(X_2) \in \mathcal{S}_z]
\end{aligned}$$

Note that for any $X \in \mathcal{D}$, the event $\mathcal{M}(X) \in \mathcal{S}$ is the union of $\mathcal{M}_1(X) = z \wedge \mathcal{M}_2(X) \in \mathcal{S}_z$ for all $z \in \mathcal{O}_1$. Also note that for any $z_1 \in \mathcal{O}_1$ and $z_2 \in \mathcal{O}_1$ that are different, the event $\mathcal{M}_1(X) =$

$z_1 \wedge \mathcal{M}_2(X) \in \mathcal{S}_{z_1}$ and the event $\mathcal{M}_1(X) = z_2 \wedge \mathcal{M}_2(X) \in \mathcal{S}_{z_2}$ are disjoint. Hence, suming up (if $\mathcal{O}_1$ is discrete) or integrating (if $\mathcal{O}_1$ is non-discrete) both sides of the above inequality for all $z \in \mathcal{O}_1$, we get that

$$\Pr[\mathcal{M}(X_1) \in \mathcal{S}] \leq \exp(\epsilon_1 + \epsilon_2) \cdot \Pr[\mathcal{M}(X_2) \in \mathcal{S}]$$

■

The following theorem gives a useful tool in designing differentially private mechanisms: designing a simple differentially private algorithm first and adding deterministic actions on top of it.

Theorem 1.4 *Let $f : \mathcal{D} \to \mathcal{O}_1$ be a ϵ-differentially private randomized algorithm, and let $g : \mathcal{O}_1 \to \mathcal{O}_2$ be a deterministic function. Then, $g \circ f : \mathcal{D} \to \mathcal{O}_2$, whose value at $X \in \mathcal{D}$ is $g(f(X))$, preserves ϵ-differential privacy.*

The proof of the Theorem 1.4 is left as an exercise.

2 Achieving Differential Privacy

Given a set $\mathcal{D}$ of databases, a deterministic function $f : \mathcal{D} \to \mathbb{R}^d$, and a privacy parameter ϵ, we want to find an ϵ-differentially private version of f, i.e., to find a (random) function $\widehat{f} : \mathcal{D} \to \mathbb{R}^d$ that has the following properties:

1. **Privacy.** The function $\widehat{f}$ preserves ϵ-differential privacy,
2. **Utility.** For any database $X \in \mathcal{D}$, with high probability, $\widehat{f}(X)$ is close to $f(X)$.

Observe that the trivial function $\widehat{f} \equiv 0$ satisfies ϵ-differential privacy. However, this will not be very useful. Here is one way to define utility.

Definition 2.1 ((λ, δ)-useful) *Let $0 < \delta < 1$ and $\lambda > 0$ be utility parameters. Let $f : \mathcal{D} \to \mathbb{R}^d$ be a deterministic function and $\widehat{f} : \mathcal{D} \to \mathbb{R}^d$ be a randomized function. We say $\widehat{f}$ is (λ, δ)-useful with respect to f, if for any database $X \in \mathcal{D}$, with probability at least $1 - \delta$, for each $i \in [d]$, $|f_i(X) - \widehat{f}_i(X)| \leq \lambda$.*

2.1 Achieving differential privacy by adding Laplace noise

One way to convert a function $f : \mathcal{D} \to \mathbb{R}^d$ into a differentially private version is to add independent random noise to each of its coordinates. Intuitively, the more $f(X)$ changes when we change one coordinate of X, the larger the random noise is needed to hide the difference. We use ℓ_1-sensitivity to formally measure the maximum difference between the values of f of any two neighboring databases.

Definition 2.2 (ℓ_1-Sensitivity) *Let $f : \mathcal{D} \to \mathbb{R}^d$ be a deterministic function. The ℓ_1-sensitivity of f, denoted by Δf, is*

$$\max_{X^{(1)} \sim X^{(2)}} ||f(X^{(1)}) - f(X^{(2)})||_1 = \max_{X^{(1)} \sim X^{(2)}} \sum_{i=1}^{d} |f_i(X^{(1)}) - f_i(X^{(2)})|$$

We use random variables sampled from Laplace distributions as the random noise.

Definition 2.3 (Laplace Distribution) *Let $b > 0$. We denote by $\mathsf{Lap}(b)$ the Laplace distribution such that the probability density function at z is $\frac{1}{2b}\exp(-\frac{|z|}{b})$.*

Laplace distribution has the following properties.

Theorem 2.4 *Let $b > 0$ and let γ be a random variable sampled from $\mathsf{Lap}(b)$. Then,*

1. $E[\gamma] = 0$,
2. $var[\gamma] = 2b^2$,
3. *for any* $\lambda > 0$, $\Pr[|\gamma| > \lambda] = \exp(-\frac{\lambda}{b})$.

Proof:

1.

$$\begin{aligned}
E[\gamma] &= \int_{-\infty}^{\infty} \frac{x}{2b}\exp\left(-\frac{|x|}{b}\right)dx = \int_{-\infty}^{0} \frac{x}{2b}\exp\left(\frac{x}{b}\right)dx + \int_{0}^{\infty} \frac{x}{2b}\exp\left(\frac{-x}{b}\right)dx \\
&= \int_{\infty}^{0} \frac{-x}{2b}\exp\left(\frac{-x}{b}\right)d(-x) + \int_{0}^{\infty} \frac{x}{2b}\exp\left(\frac{-x}{b}\right)dx \\
&= -\int_{0}^{\infty} \frac{x}{2b}\exp\left(\frac{-x}{b}\right)dx + \int_{0}^{\infty} \frac{x}{2b}\exp\left(\frac{-x}{b}\right)dx = 0
\end{aligned}$$

2.

$$\begin{aligned}
E[\gamma^2] &= \int_{-\infty}^{\infty} \frac{x^2}{2b}\exp\left(-\frac{|x|}{b}\right)dx = 2\int_{0}^{\infty} \frac{x^2}{2b}\exp\left(-\frac{|x|}{b}\right)dx \\
&= -\int_{0}^{\infty} x^2\exp\left(-\frac{x}{b}\right)d\left(-\frac{x}{b}\right) = -(x^2\exp\left(-\frac{x}{b}\right)|_0^{\infty} - \int_{0}^{\infty} 2x\exp\left(-\frac{x}{b}\right)dx) \\
&= 2\int_{0}^{\infty} x\exp\left(-\frac{x}{b}\right)dx = -2b\int_{0}^{\infty} x\exp\left(-\frac{x}{b}\right)d\left(-\frac{x}{b}\right) \\
&= -2b\left(x\exp\left(-\frac{x}{b}\right)|_0^{\infty} - \int_{0}^{\infty}\exp\left(-\frac{x}{b}\right)dx\right) = 2b\int_{0}^{\infty}\exp\left(-\frac{x}{b}\right)dx \\
&= -2b^2\int_{0}^{\infty}\exp\left(-\frac{x}{b}\right)d\left(-\frac{x}{b}\right) = -2b^2\exp\left(-\frac{x}{b}\right)\Big|_0^{\infty} = 2b^2
\end{aligned}$$

Hence,

$$var[\gamma] = E[\gamma^2] - (E[\gamma])^2 = 2b^2 - 0 = 2b^2$$

3.

$$\begin{aligned}\Pr[|\gamma| > \lambda] &= \Pr[\gamma > \lambda] + \Pr[\gamma < -\lambda] = \int_{\lambda}^{\infty} \frac{1}{2b} \exp\left(-\frac{x}{b}\right) dx + \int_{-\infty}^{-\lambda} \frac{1}{2b} \exp\left(\frac{x}{b}\right) dx \\ &= 2\int_{\lambda}^{\infty} \frac{1}{2b} \exp\left(-\frac{x}{b}\right) dx = -\int_{\lambda}^{\infty} \exp\left(-\frac{x}{b}\right) d\left(-\frac{x}{b}\right) = -\exp\left(-\frac{x}{b}\right)\Big|_{\lambda}^{\infty} \\ &= \exp\left(-\frac{\lambda}{b}\right)\end{aligned}$$

■

Note that if b is small, the mass is highly concentrated around 0. Hence, it can only be used to privatize functions with small sensitivity. However, the highly concentrated mass also implies good utility. The following theorem shows that choosing $b := \frac{\Delta f}{\epsilon}$ is enough to preserve privacy and with high probability, the additive error incurred by the random noise is small.

Theorem 2.5 *Let $f : \mathcal{D} \to \mathbb{R}^d$ be a deterministic function, $0 < \epsilon < 1$ be the privacy parameter and $0 < \delta < 1$ be the failure probability. Let $\gamma_1, \gamma_2, \ldots, \gamma_d$ be random variables independently sampled from $\mathsf{Lap}(\frac{\Delta f}{\epsilon})$. Then, the randomized function $\widehat{f}$ such that $\widehat{f}_i(X) := f_i(x) + \gamma_i$ for all $i \in [d]$*

1. *preserves ϵ-differential privacy,*
2. *is $(\frac{\Delta f}{\epsilon} \ln \frac{d}{\delta}, \delta)$-useful with respect to f.*

Proof: Let $X^{(1)} \in \mathcal{D}$ and $X^{(2)} \in \mathcal{D}$ be two neighboring database. Let $z \in \mathbb{R}^d$ be a vector. We abuse the notation a little bit and use $\Pr[\widehat{f}(X^{(1)}) = z]$ and $\Pr[\widehat{f}(X^{(2)}) = z]$ to denote the density instead of the probability. Hence, we have

$$\begin{aligned}\frac{\Pr[\widehat{f}(X^{(1)}) = z]}{\Pr[\widehat{f}(X^{(2)}) = z]} &= \frac{\Pr[\wedge_{i=1}^{d} f_i(X^{(1)}) + \gamma_i = z_i]}{\Pr[\wedge_{i=1}^{d} f_i(X^{(2)}) + \gamma_i = z_i]} \\ &= \frac{\Pr[\wedge_{i=1}^{d} \gamma_i = z_i - f_i(X^{(1)})]}{\Pr[\wedge_{i=1}^{d} \gamma_i = z_i - f_i(X^{(2)})]} \\ &= \frac{\prod_{i=1}^{d} \Pr[\gamma_i = z_i - f_i(X^{(1)})]}{\prod_{i=1}^{d} \Pr[\gamma_i = z_i - f_i(X^{(2)})]} \\ &= \frac{\prod_{i=1}^{d} \frac{\epsilon}{2\Delta f} \exp(-\frac{|f_i(X^{(1)}) - z_i|}{\Delta f/\epsilon})}{\prod_{i=1}^{d} \frac{\epsilon}{2\Delta f} \exp(-\frac{|f_i(X^{(2)}) - z_i|}{\Delta f/\epsilon})} \\ &= \exp\left(\sum_{i=1}^{d} \left(\frac{|f(X^{(2)}) - z_i|}{\Delta f/\epsilon} - \frac{|f(X^{(1)}) - z_i|}{\Delta f/e}\right)\right) \\ &\leq \exp\left(\sum_{i=1}^{d} \frac{|f(X^{(1)}) - f(X^{(2)})|}{\Delta f/\epsilon}\right) \\ &\leq \exp\left(\frac{\Delta f}{\Delta f/\epsilon}\right) \\ &= \exp(\epsilon)\end{aligned}$$

For any measurable subset $S \subseteq \mathbb{R}^d$, $\Pr[\widehat{f}(X^{(1)}) \in S] = \int_S \Pr[\widehat{f}(X^{(1)}) = z]dz \leq \exp(\epsilon) \int_S \Pr[\widehat{f}(X^{(2)}) = z]dz = \Pr[\widehat{f}(X^{(2)}) \in S]$. Hence, the privacy guarantee is proved.

By Property 3 of Theorem 2.4, we know that for all $i \in [d]$, $\Pr[|\gamma_i| > \frac{\Delta f}{\epsilon} \ln \frac{d}{\delta}] = \exp(-(\frac{\Delta f}{\epsilon} \ln \frac{d}{\delta})/(\frac{\Delta f}{\epsilon})) = \frac{\delta}{d}$. Hence, by union bound on $i \in [d]$, we know that $\Pr[\vee_{i\in[d]}|\gamma_i| > \frac{\Delta f}{\epsilon} \ln \frac{d}{\delta}] \leq \delta$.

Let $X \in \mathcal{D}$ be any database. Note that $|f_i(X) - \widehat{f}_i(X)| = |\gamma_i|$ for all $i \in [d]$. Hence, by the union bound, $\Pr[\exists i \in [d], |f_i(X) - \widehat{f}_i(X)| > \frac{\Delta f}{\epsilon} \ln \frac{d}{\delta}] \leq \delta$, which is equivalent to $\Pr[\wedge_{i\in[d]}|f_i(X) - \widehat{f}_i(X)| \leq \frac{\Delta f}{\epsilon} \ln \frac{d}{\delta}] \geq 1 - \delta$. Thus, the utility guarantee is proved. ■

2.2 Achieving differential privacy by adding geometric noise

If the function f has only integer values, we can add the discrete counterpart of Laplace noise, geometric noise, to achieve differential privacy.

Definition 2.6 (Geometric Distribution) *Let $\alpha > 1$. We denote by $\mathsf{Geom}(\alpha)$ the symmetric geometric distribution that takes integer values such that the probability mass function at k is $\frac{\alpha-1}{\alpha+1} \cdot \alpha^{-|k|}$.*

Geometric distribution have similar properties to those of Laplace distribution.

Theorem 2.7 *Let $\alpha > 1$, and let γ be a random variable sampled from symmetric geometric distribution $\mathsf{Geom}(\alpha)$. Then,*

1. *$E[\gamma] = 0$,*
2. *$var[\gamma] = \frac{2\alpha}{(1-\alpha)^2}$,*
3. *for any integer $z \geq 0$, $\Pr[|\gamma| > z] \leq \frac{1}{\alpha^z}$.*

Similar to Theorem 2.5, we have the following result:

Theorem 2.8 *Let $f : \mathcal{D} \to \mathbb{Z}^d$ be a deterministic function, $0 < \epsilon < 1$ be the privacy parameter and $0 < \delta < 1$ be the failure probability. Let $\gamma_1, \gamma_2, \ldots, \gamma_d$ be random variables independently sampled from $\mathsf{Geom}(\exp(\frac{\epsilon}{\Delta f}))$. Then, the randomized function $\widehat{f}$ such that $\widehat{f}_i(X) := f_i(x) + \gamma_i$ for all $i \in [d]$*

1. *preserves ϵ-differential privacy,*
2. *is $(\frac{\Delta f}{\epsilon} \ln \frac{d}{\delta}, \delta)$-useful with respect to f.*

The proof of the above two theorems are left as exercises.

3 Homework Preview

1. Prove Theorem 2.7.
2. Prove Theorem 2.8.
3. Prove Theorem 1.4.