Answering Numeric Queries

The Laplace Distribution:

Lap(b) is the probability distribution with p.d.f.:

$$p(x \mid b) = \frac{1}{2b}\exp\left(-\frac{|x|}{b}\right)$$

i.e. a symmetric exponential distribution

$Y \sim \mathrm{Lap}(b), \qquad E[|Y|] = b$

$\Pr[|Y| \geq t \cdot b] = e^{-t}$

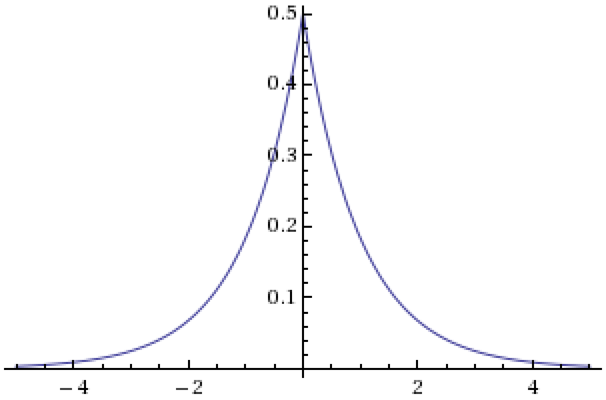

Answering Numeric Queries: The Laplace Mechanism

Laplace$(D, Q\colon \mathbb{N}^{|X|} \to \mathbb{R}^k, \epsilon)$:

1. Let $\Delta = GS(Q)$.
2. For $i = 1$ to k: Let $Y_i \sim \mathrm{Lap}(\frac{\Delta}{\epsilon})$.
3. Output $Q(D) + (Y_1, \ldots, Y_k)$

Independently perturb each coordinate of the output with Laplace noise scaled to the sensitivity of the function.

Idea: This should be enough noise to hide the contribution of any single individual, no matter what the database was.

Answering Numeric Queries: The Laplace Mechanism

Laplace($D, Q: \mathbb{N}^{|X|} \rightarrow \mathbb{R}^k, \epsilon$):

1. Let $\Delta = GS(Q)$.
2. For $i = 1$ to k: Let $Y_i \sim \mathrm{Lap}(\frac{\Delta}{\epsilon})$.
3. Output $Q(D) + (Y_1, \ldots, Y_k)$

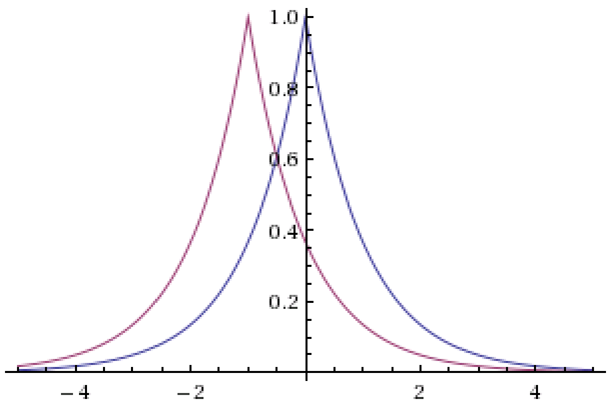

Answering Numeric Queries: The Laplace Mechanism

Theorem: The Laplace mechanism is $(\epsilon, 0)$-differentially private.

Proof:

Consider any pair of databases D, D' with $\left||D - D'|\right|_1 \leq 1$

Consider any event $S \subseteq \mathbb{R}^k$

$$\frac{\Pr[\text{Laplace}(D, Q, \epsilon) \in S]}{\Pr[\text{Laplace}(D', Q, \epsilon) \in S]} = \frac{\int_{x \in S} \Pr[\text{Laplace}(D, Q, \epsilon) = x]}{\int_{x \in S} \Pr[\text{Laplace}(D', Q, \epsilon) = x]}$$

$$\leq \max_{x \in S} \frac{\Pr[\text{Laplace}(D, Q, \epsilon) = x]}{\Pr[\text{Laplace}(D', Q, \epsilon) = x]}$$

Answering Numeric Queries: The Laplace Mechanism

Theorem: The Laplace mechanism is $(\epsilon, 0)$-differentially private.

Proof: Let $y = \text{Laplace}(D, Q, \epsilon), y' = \text{Laplace}(D', Q, \epsilon)$

$$\frac{\Pr[y = x]}{\Pr[y' = x]} = \prod_{i=1}^{k} \frac{\Pr[y_i = x_i]}{\Pr[y_i' = x_i]} = \prod_{i=1}^{k} \frac{\Pr[Q(D)_i + Y_i = x_i]}{\Pr[Q(D')_i + Y_i = x_i]}$$

$$= \prod_{i=1}^{k} \frac{\Pr[Y_i = x_i - Q(D)_i]}{\Pr[Y_i = x_i - Q(D')_i]} = \prod_{i=1}^{k} \frac{\exp(-\epsilon \frac{|x_i - Q(D)_i|}{\Delta})}{\exp(-\epsilon \frac{|x_i - Q(D')_i|}{\Delta})}$$

$$= \prod_{i=1}^{k} \exp\left(\epsilon \frac{|x_i - Q(D')_i| - |x_i - Q(D)_i|}{\Delta}\right) \leq \prod_{i=1}^{k} \exp\left(\epsilon \frac{|Q(D)_i - Q(D')_i|}{\Delta}\right)$$

$$= \exp\left(\frac{\epsilon}{\Delta} \sum_{i=1}^{k} |Q(D)_i - Q(D')_i|\right) \leq \exp\left(\frac{\epsilon}{\Delta} \Delta\right) = \exp(\epsilon)\,.$$

Answering Numeric Queries: The Laplace Mechanism

Take away message:

1) Low sensitivity queries can be answered with very little noise!
   $\mathrm{E}\left[\mathrm{Lap}\left(\frac{1}{\epsilon}\right)\right] = \frac{1}{\epsilon}$
2) Comparison with lower bound for blatant non-privacy: A subset-sum query $Q: \{0,1\}^{|X|} \rightarrow \mathbb{R}$ has sensitivity $\mathrm{GS}(Q) = 1$. Any k of them jointly have sensitivity k. So Laplace Mechanism lets you answer any k subset-sum queries with error $o(\frac{k}{\epsilon})$
   1) Recall: Lower bound for blatant non-privacy required $2^{|X|}$ subset-sum queries with error $o(|X|)$ or $\tilde{O}(|X|)$ queries with error $o(\sqrt{|X|})$. So there is a gap we can close!

Privacy for Non-Numeric Queries

The Exponential Mechanism

Output Perturbation

- We know how to handle (a single) numeric query.
  - “How many people in this room have blue eyes?”
  - Perturb the answer by an amount proportional to the sensitivity of the query.
    - Noise of magnitude $O\left(\frac{\Delta}{\epsilon}\right)$ drawn from the Laplace distribution suffices for $(\epsilon, 0)$-differential privacy

When Output Perturbation Doesn’t Make Sense

- What about if we have a non-numeric valued query?
  - “What is the most common eye color in this room?”
- What if the perturbed answer isn’t almost as good as the exact answer?
  - “Which price would bring the most money from a set of buyers?

Example: Items for sale

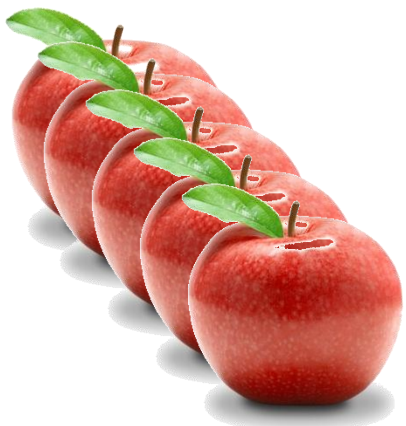

Could set the price of apples at $1.00 for profit: $4.00

Could set the price of apples at $4.01 for profit $4.01

Best price: $4.01
2nd best price: $1.00
Profit if you set the price at $4.02: $0
Profit if you set the price at $1.01: $1.01

$1.00

$1.00

$1.00

$4.01

The Exponential Mechanism

- A mechanism $M: \mathbb{N}^{|X|} \rightarrow R$ for some abstract range R.
  - i.e. $R = \{\text{Red, Blue, Green, Brown, Purple}\}$
  - $R = \{\$1.00, \$1.01, \$1.02, \$1.03, \dots\}$
- Paired with a *quality score:*

$$q: \mathbb{N}^{|X|} \times R \rightarrow \mathbb{R}$$

$q(D, r)$ represents how good output r is for database D.

The Exponential Mechanism

- Relative parameters for privacy, solution quality:
  - Sensitivity of q:
    $$GS(q) = \max_{r \in R, D, D' : \left\| D - D' \right\|_1 \leq 1} |q(D,r) - q(D',r)|$$
  - Size and structure of R.
    - How many elements of R are high quality? How many are low quality?

The Exponential Mechanism

Exponential$(D, R, q: \mathbb{N}^{|X|} \to R, \epsilon)$:

1. Let $\Delta = GS(q)$.
2. Output $r \sim R$ with probability proportional to:

$$\Pr[r] \sim \exp\left(\frac{\epsilon q(D,r)}{2\Delta}\right)$$

$$\Pr[r] = \frac{\exp(\frac{\epsilon q(D,r)}{2\Delta})}{\sum_{r' \in R} \exp(\frac{\epsilon q(D,r')}{2\Delta})}$$

The Exponential Mechanism

Exponential$(D, R, q: \mathbb{N}^{|X|} \rightarrow R, \epsilon)$:

1. Let $\Delta = GS(q)$.
2. Output $r \sim R$ with probability proportional to:

$$\Pr[r] \sim \exp\left(\frac{\epsilon q(D,r)}{2\Delta}\right)$$

Idea: Make high quality outputs exponentially more likely at a rate that depends on the sensitivity of the quality score (and the privacy parameter)

The Exponential Mechanism

Exponential$(D, R, q: \mathbb{N}^{|X|} \to R, \epsilon)$:

1. Let $\Delta = GS(q)$.
2. Output $r \sim R$ with probability proportional to:

$$\Pr[r] \sim \exp\left(\frac{\epsilon q(D, r)}{2\Delta}\right)$$

Theorem: The Exponential Mechanism preserves $(\epsilon, 0)$-differential privacy.

The Exponential Mechanism

Exponential$(D, R, q\colon \mathbb{N}^{|X|} \rightarrow R, \epsilon)$:

1. Let $\Delta = GS(q)$.
2. Output $r \sim R$ with probability proportional to:

$$\Pr[r] \sim \exp\left(\frac{\epsilon q(D,r)}{2\Delta}\right)$$

Theorem: The Exponential Mechanism preserves $(\epsilon, 0)$-differential privacy.

Proof: Fix any $D, D' \in \mathbb{N}^{|X|}$ with $\left|\left|D, D'\right|\right|_1 \leq 1$ and any $r \in R$…

$$\frac{\Pr[\text{Exponential}(D,R,q,\epsilon)=r]}{\Pr[\text{Exponential}(D',R,q,\epsilon)=r]} =$$

$$\frac{\left(\dfrac{\exp(\frac{\epsilon q(D,r)}{2\Delta})}{\sum \exp(\frac{\epsilon q(D,r')}{2\Delta})}\right)}{\left(\dfrac{\exp(\frac{\epsilon q(D',r)}{2\Delta})}{\sum \exp(\frac{\epsilon q(D',r')}{2\Delta})}\right)} = \left(\frac{\exp(\frac{\epsilon q(D,r)}{2\Delta})}{\exp(\frac{\epsilon q(D',r)}{2\Delta})}\right)\left(\frac{\sum_{r'} \exp(\frac{\epsilon q(D',r')}{2\Delta})}{\sum_{r'} \exp(\frac{\epsilon q(D,r')}{2\Delta})}\right)$$

The Exponential Mechanism

Exponential$(D, R, q\colon \mathbb{N}^{|X|} \rightarrow R, \epsilon)$:

1. Let $\Delta = GS(q)$.
2. Output $r \sim R$ with probability proportional to:

$$\Pr[r] \sim \exp\left(\frac{\epsilon q(D, r)}{2\Delta}\right)$$

Theorem: The Exponential Mechanism preserves $(\epsilon, 0)$-differential privacy.

Proof:

$$\bigstar \quad = \left(\frac{\exp(\frac{\epsilon q(D, r)}{2\Delta})}{\exp(\frac{\epsilon q(D', r)}{2\Delta})}\right) =$$

$$\exp\left(\frac{\epsilon(q(D, r) - q(D', r))}{2\Delta}\right) \leq$$

$$\exp\left(\frac{\epsilon\Delta}{2\Delta}\right) = \exp\left(\frac{\epsilon}{2}\right)$$

The Exponential Mechanism

Exponential$(D, R, q\colon \mathbb{N}^{|X|} \to R, \epsilon)$:
1. Let $\Delta = GS(q)$.
2. Output $r \sim R$ with probability proportional to:

$$\Pr[r] \sim \exp\left(\frac{\epsilon q(D,r)}{2\Delta}\right)$$

Theorem: The Exponential Mechanism preserves $(\epsilon, 0)$-differential privacy.

Proof:

$$\star\star = \left(\frac{\sum_{r'} \exp(\frac{\epsilon q(D',r')}{2\Delta})}{\sum_{r'} \exp(\frac{\epsilon q(D,r')}{2\Delta})}\right) \leq$$

$$\left(\frac{\sum_{r'} \exp(\frac{\epsilon (q(D,r') + \Delta)}{2\Delta})}{\sum_{r'} \exp(\frac{\epsilon q(D,r')}{2\Delta})}\right) =$$

$$= \left(\frac{\exp(\frac{\epsilon}{2}) \sum_{r'} \exp(\frac{\epsilon q(D,r')}{2\Delta})}{\sum_{r'} \exp(\frac{\epsilon q(D,r')}{2\Delta})}\right) = \exp(\frac{\epsilon}{2})$$

The Exponential Mechanism

Exponential$(D, R, q\colon \mathbb{N}^{|X|} \to R, \epsilon)$:
1. Let $\Delta = GS(q)$.
2. Output $r \sim R$ with probability proportional to:

$$\Pr[r] \sim \exp\left(\frac{\epsilon q(D, r)}{2\Delta}\right)$$

Theorem: The Exponential Mechanism preserves $(\epsilon, 0)$-differential privacy.

Proof: Recall:

$$\frac{\Pr[\mathrm{Exponential}(D,R,q,\epsilon)=r]}{\Pr[\mathrm{Exponential}(D\prime,R,q,\epsilon)=r]} =$$

$$\leq \exp\left(\frac{\epsilon}{2}\right) \exp\left(\frac{\epsilon}{2}\right)$$

$$= \exp(\epsilon)$$

The Exponential Mechanism

Exponential$(D, R, q: \mathbb{N}^{|X|} \to R, \epsilon)$:

1. Let $\Delta = GS(q)$.
2. Output $r \sim R$ with probability proportional to:

$$\Pr[r] \sim \exp\left(\frac{\epsilon q(D,r)}{2\Delta}\right)$$

But is the answer any good?

The Exponential Mechanism

Exponential$(D, R, q: \mathbb{N}^{|X|} \rightarrow R, \epsilon)$:

1. Let $\Delta = GS(q)$.
2. Output $r \sim R$ with probability proportional to:

$$\Pr[r] \sim \exp\left(\frac{\epsilon q(D, r)}{2\Delta}\right)$$

But is the answer any good?

It depends…

The Exponential Mechanism

Define:

$$OPT_q(D) = \max_{r \in R} q(D, r)$$
$$R_{OPT} = \{r \in R : q(D, r) = OPT_q(D)\}$$
$$r^* = \text{Exponential}(D, R, q, \epsilon)$$

Theorem:

$$\Pr\left[q(r^*) \leq OPT_q(D) - \frac{2\Delta}{\epsilon}\left(\log\left(\frac{|R|}{|R_{OPT}|}\right) + t\right)\right] \leq e^{-t}$$

The Exponential Mechanism

Theorem:

$$\Pr\left[q(r^*) \leq OPT_q(D) - \frac{2\Delta}{\epsilon}\left(\log\left(\frac{|R|}{|R_{OPT}|}\right) + t\right)\right] \leq e^{-t}$$

Corollary:

$$\Pr\left[q(r^*) \leq OPT_q(D) - \frac{2\Delta}{\epsilon}(\log(|R|) + t)\right] \leq e^{-t}$$

Proof:

$|R_{OPT}| \geq 1$ by definition.

The Exponential Mechanism

Theorem:

$$\Pr\left[q(r^*) \leq OPT_q(D) - \frac{2\Delta}{\epsilon}\left(\log\left(\frac{|R|}{|R_{OPT}|}\right) + t\right)\right] \leq e^{-t}$$

Corollary:

$$\mathrm{E}[q(r^*)] \geq OPT_q(D) \ -\frac{2\Delta}{\epsilon}\left(\log(|R|) + \log\left(OPT_q(D)\right)\right) - 1$$

Proof:

$$\Pr\left[q(r^*) \leq OPT_q(D) - \frac{2\Delta}{\epsilon}\left(\log(|R|) + \log(OPT_q(D)\right)\right] \leq \frac{1}{OPT_q(D)}$$

$$\Pr\left[q(r^*) \geq OPT_q(D) - \frac{2\Delta}{\epsilon}\left(\log(|R|) + \log(OPT_q(D)\right)\right] \geq 1 - \frac{1}{OPT_q(D)}$$

The Exponential Mechanism

Theorem:

$$\Pr\left[q(r^*) \leq OPT_q(D) - \frac{2\Delta}{\epsilon}\left(\log\left(\frac{|R|}{|R_{OPT}|}\right) + t\right)\right] \leq e^{-t}$$

Corollary:

$$\mathrm{E}[q(r^*)] \geq OPT_q(D) \; - \frac{2\Delta}{\epsilon}\left(\log(|R|) + \log\left(OPT_q(D)\right)\right) - 1$$

Proof:

$$E[q(r^*)] \geq (x \cdot \Pr[q(r^*) \geq x])$$

$$\geq \left(OPT_q(D) \; - \frac{2\Delta}{\epsilon}\left(\log(|R|) + \log\left(OPT_q(D)\right)\right)\right) \cdot \left(1 - \frac{1}{OPT_q(D)}\right)$$

$$> OPT_q(D) \; - \frac{2\Delta}{\epsilon}\left(\log(|R|) + \log\left(OPT_q(D)\right)\right) - 1$$

The Exponential Mechanism

Theorem:

$$\Pr\left[q(r^*) \leq OPT_q(D) - \frac{2\Delta}{\epsilon}\left(\log\left(\frac{|R|}{|R_{OPT}|}\right) + t\right)\right] \leq e^{-t}$$

Proof:

$$\Pr[q(r^*) \leq x] \leq \frac{\Pr[q(r^*) \leq x]}{\Pr\left[q(r^*) = OPT_q(D)\right]}$$

$$\leq \frac{|\mathrm{R}|\exp(\frac{\epsilon \mathrm{x}}{2\Delta})}{|\mathrm{R_{OPT}}|\exp(\frac{\epsilon OPT_q(D)}{2\Delta})}$$

$$= \frac{|R|}{|R_{OPT}|}\exp\left(\frac{\epsilon\left(x - OPT_q(D)\right)}{2\Delta}\right) = \left(\frac{|R|}{|R_{OPT}|}\right)\exp\left(-\log\left(\frac{|R|}{|R_{OPT}|}\right) - t\right)$$

$$= \left(\frac{|R|}{|R_{OPT}|}\right)\left(\frac{|R_{OPT}|}{|R|}\right)e^{-t} = e^{-t}$$

The Exponential Mechanism

So if $R = \{\text{Red, Blue, Green, Brown, Purple}\}$ then we can answer "What is the most common eye color in this room?" with a color that is shared by:

$OPT - \frac{2}{\epsilon}(\log(5) + 3) < OPT - \frac{7.4}{\epsilon}$ people

Except with probability: $\leq e^{-3} < .05$

Independent of the number of people in the room.

Very small error if n is large.

To Muse On

- The exponential mechanism is based on the vector: $\hat{q}: \mathbb{N}^{|X|} \rightarrow |R| = \left(q(D, r_1), q(D, r_2), \ldots, q\left(D, r_{|R|}\right)\right)$
  - Might have sensitivity $GS(\hat{q}) = |R| \cdot GS(q)$.
  - Exponential Mechanism only depends on $GS(q)$.
- Error has only logarithmic dependence on $|R|$.
  - Could take exponentially large ranges!
  - But *sampling* from the exponential mechanism efficiently is non-trivial.

To Muse On

- Read [MT07]: “Mechanism Design Via Differential Privacy”
  - Introduces the exponential mechanism
  - Blog post: Describe the application of the exponential mechanism to digital goods auctions.