Algorithms for Differential Privacy: Exponential & Median Mechanism

CompSci 590.03

Instructor: Ashwin Machanavajjhala

Duke University

Recap: Differential Privacy

For every pair of inputs that differ in one value

For every output …

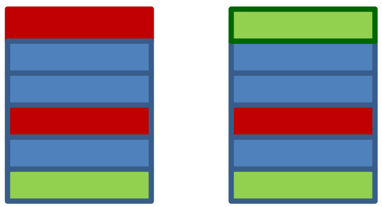

D_1 *D_2*

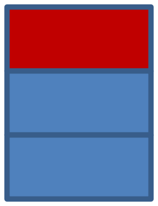

O

Adversary should not be able to distinguish between any D_1 and D_2 based on any O

$$\log\left(\frac{\Pr[A(D_1) = O]}{\Pr[A(D_2) = O]}\right) < \varepsilon \qquad (\varepsilon > 0)$$

Recap: Differential Privacy

- For every pair of tables D1 and D2,
  adversary should not be able to distinguish between D1 and D2.

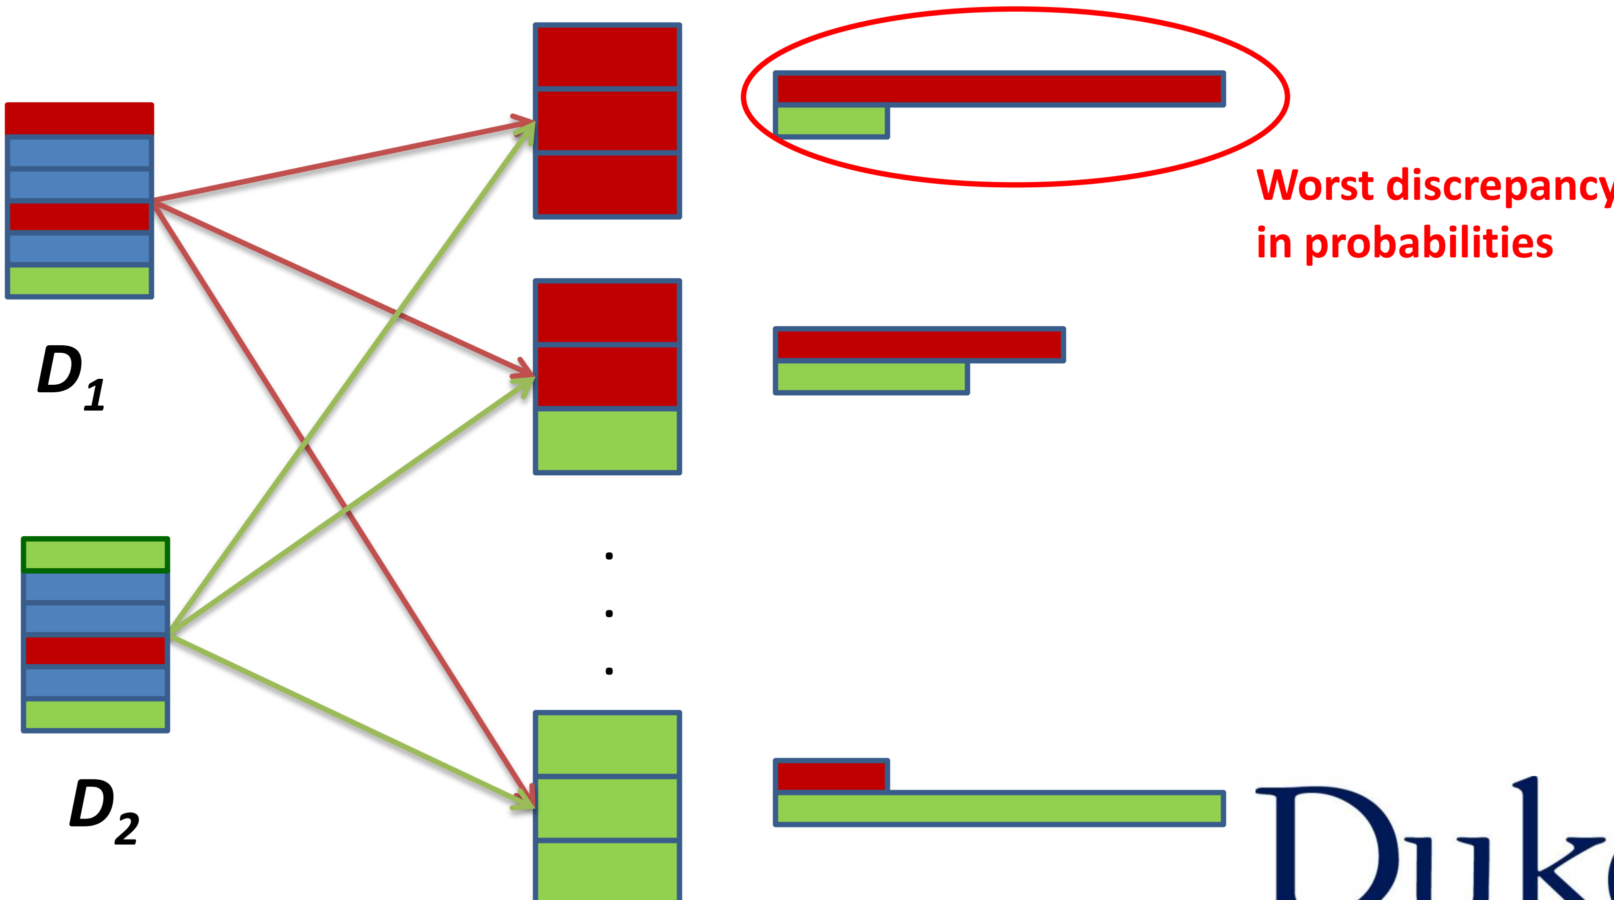

Composability of Differential Privacy

Theorem **(Composability)**:

If algorithms $A_1, A_2, \ldots, A_k$ use independent randomness and each A_i satisfies $\boldsymbol{\varepsilon_i}$-differential privacy, resp.

Then, outputting all the answers together satisfies differential privacy with

$$\boldsymbol{\varepsilon = \varepsilon_1 + \varepsilon_2 + \ldots + \varepsilon_k}$$

Duke UNIVERSITY

Recap: Algorithms

- No deterministic algorithm guarantees differential privacy.
- Random sampling does not guarantee differential privacy.
- Randomized response satisfies differential privacy.

$$\frac{P(D \to O)}{P(D' \to O)} \leq e^{\varepsilon} \Leftrightarrow \frac{1}{1+e^{\varepsilon}} < p < \frac{e^{\varepsilon}}{1+e^{\varepsilon}}$$

Duke UNIVERSITY

Recap: Laplacian Distribution

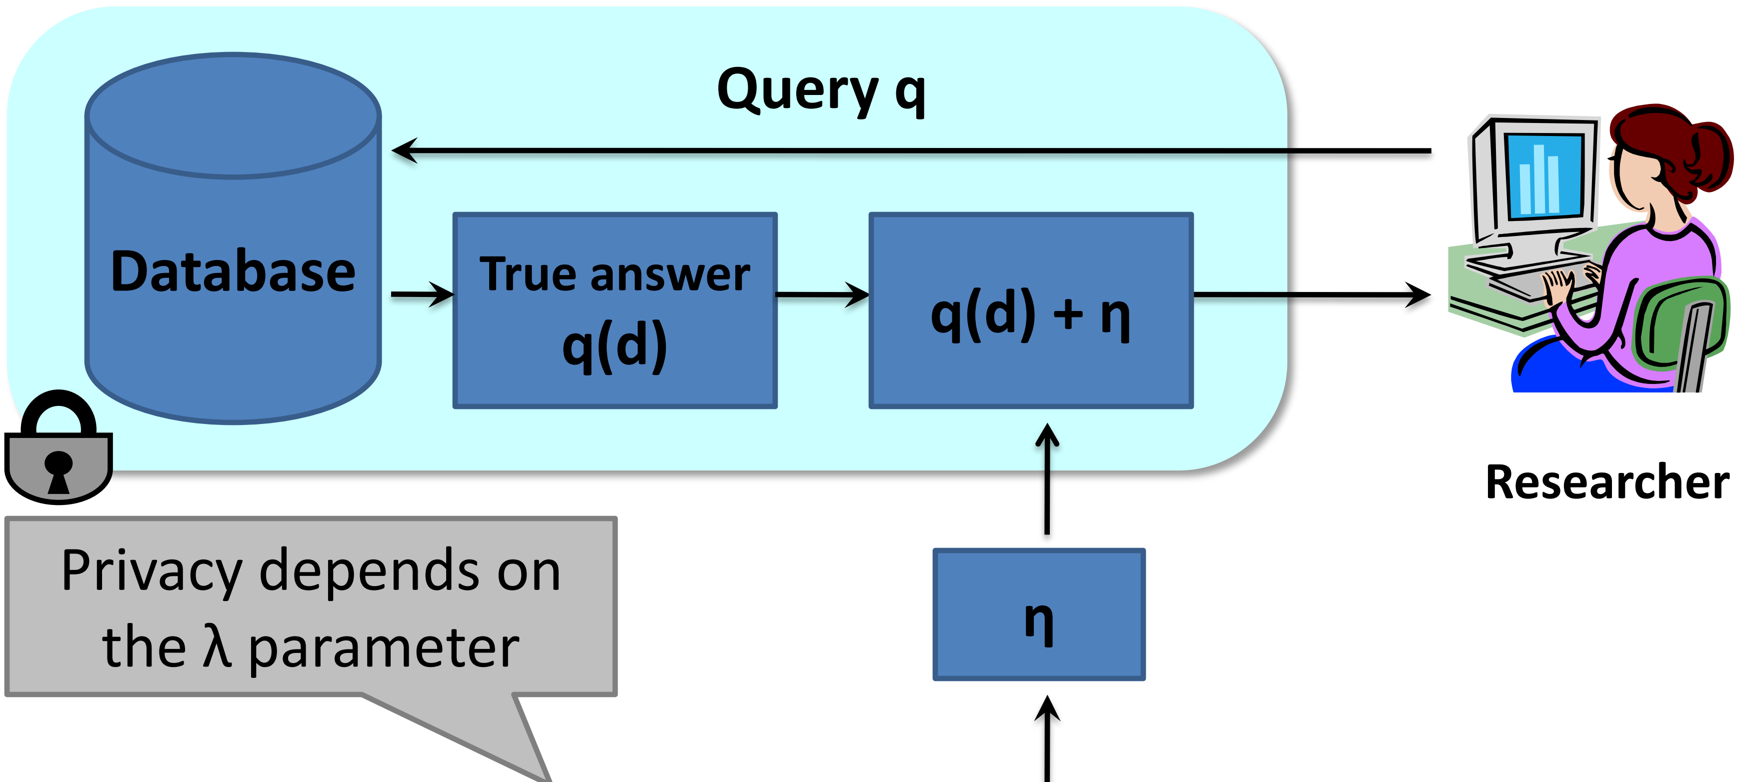

$$h(\eta) \propto \exp(-\eta / \lambda)$$

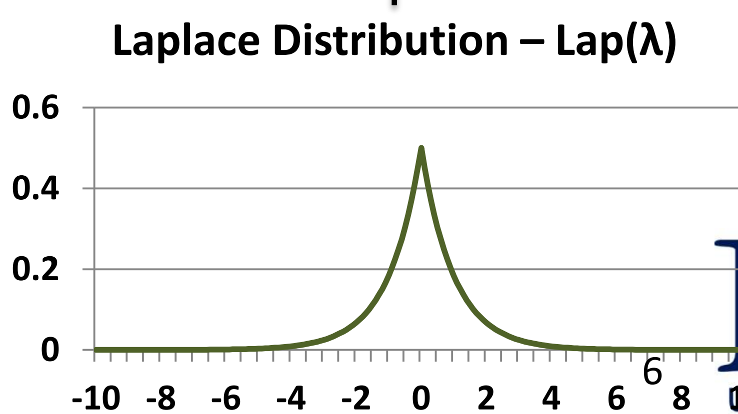

Mean: 0,
Variance: $2\lambda^2$

Duke UNIVERSITY

Recap: Laplace Mechanism

[Dwork et al., TCC 2006]

Thm: If **sensitivity** of the query is **S**, then the following guarantees ε-differential privacy.

$$\lambda = S/\varepsilon$$

Recap: Sensitivity of a Query – S(q)

[Dwork et al., TCC 2006]

Smallest number s.t. for any d, d' differing in one entry,

$$|| q(d) - q(d') || \leq S(q)$$

Example 2: **HISTOGRAM queries**

- Suppose each entry in d takes values in $\{c_1, c_2, \ldots, c_n\}$.
- Histogram(d) = $\{m_1, \ldots, m_n\}$, where m_i = (# entries in d with value c_i)
- S(q) = 2 for Histogram(d).

Changing one entry in d from c_i to c_j

- reduces the count of m_i by 1, and
- increases the count of m_j by 1.

Duke UNIVERSITY

This class

- Exponential Mechanism: when the answer is not a real number

- Median Mechanism: Answering a stream of queries

Limitations of output perturbation

- What if the answer is non-numeric?
  - "what is the most common nationality in this room": Chinese/Indian/American…
  - Other examples?

- What if the perturbed answer is not as good as the real answer?
  - "Which price would bring the most money from a set of buyers?"

Example: Items for sale

$100

$100

$100

$401

- If price is set at $100, make a revenue of $400
- If price is set at $401, make a revenue of $401

- Best price: $401, Next best: $100

- Revenue at $402 = $0
- Revenue at $101 = $101

Exponential Mechanism

- Consider some algorithm A (can be deterministic or probabilistic):

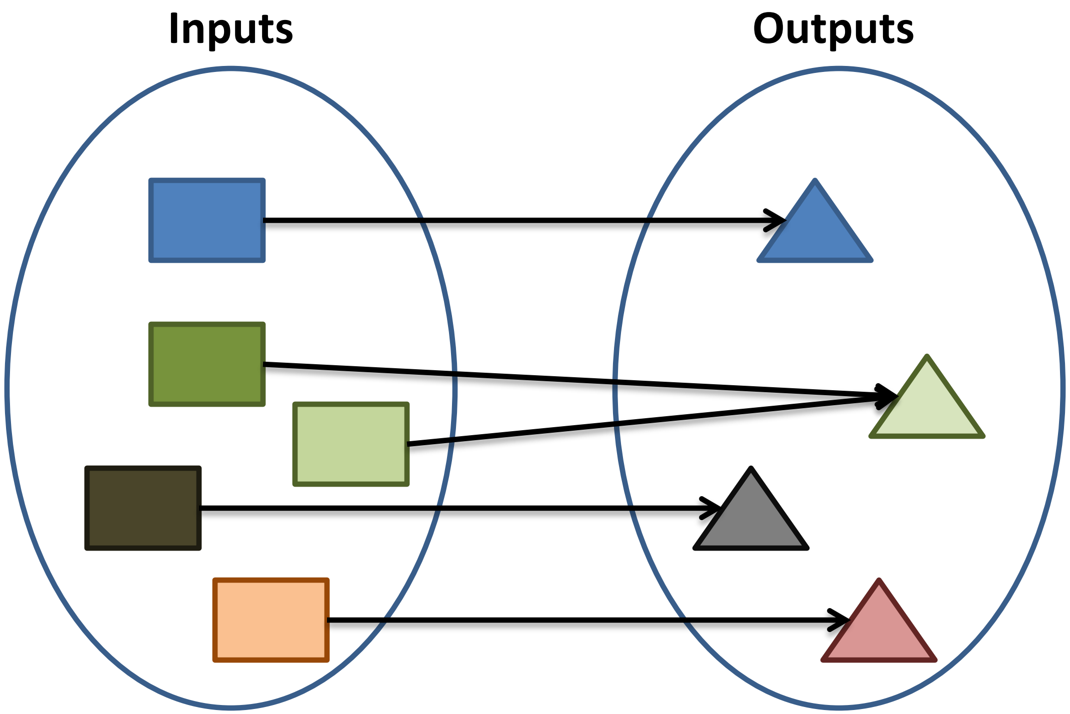

- **How to construct a differentially private version of A?**

Exponential Mechanism

- Construct a scoring function *w: Inputs x Outputs* → *R*

Examples:

- w(D, O) = c, for all D ε Inputs and O ε Outputs.
- w(D,O) = P[A(D) = O], for all D ε Inputs and O ε Outputs.

- For good utility w(D,O) should mirror the true algorithm as well as possible.

Exponential Mechanism

- Construct a scoring function *w: Inputs x Outputs* → *R*

- Sensitivity of w

$$\Delta_w = \max_{O\&D,D'} |w(D,O) - w(D,O')|$$

where D, D’ differ in one tuple

Exponential Mechanism

- Construct a scoring function *w: Inputs x Outputs* → *R*

$$Algorithm\ \mathcal{E}_w^{\varepsilon}(D)$$

- Given an input D,
  Randomly sample an output O from *Outputs* with probability

$$\frac{e^{\frac{\varepsilon}{2\Delta}\cdot w(D,O)}}{\sum_{Q \in Outputs} e^{\frac{\varepsilon}{2\Delta}\cdot w(D,Q)}}$$

Theorem

$Algorithm\ \mathcal{E}_w^{\varepsilon}(D)\ satisfies\ \varepsilon\ differential\ privacy.$

Utility of the Exponential Mechanism

- Depends on the choice of scoring function – weight given to the best output.

- E.g.,
  "What is the most common nationality?"
  w(D,nationality) = # people in D having that nationality

  Sensitivity of w is 1.

- Q: What will the output look like?

Duke UNIVERSITY

Utility of Exponential Mechanism

- Let OPT(D) = nationality with the max score
- Let O_{OPT} = {O ε Outputs : w(D,O) = OPT(D)}

- Let the exponential mechanism return an output O*

Theorem:

$$\Pr\left[w(D, O^*) \leq OPT(D) - \frac{2\Delta}{\varepsilon}\left(\log\frac{|Outputs|}{|O_{OPT}|} + t\right)\right] \leq e^{-t}$$

Utility of Exponential Mechanism

Theorem:

$$\Pr\left[w(D, O^*) \leq OPT(D) - \frac{2\Delta}{\varepsilon}\left(\log\frac{|Outputs|}{|O_{OPT}|} + t\right)\right] \leq e^{-t}$$

Suppose there are 4 nationalities

Outputs = {Chinese, Indian, American, Greek}

Exponential mechanism will output some nationality that is shared by at least K people with probability $1\text{-}e^{-3}$(=0.95), where

$$K \geq OPT - 2(\log(4) + 3)/\varepsilon = OPT - 6.8/\varepsilon$$

Duke UNIVERSITY

Laplace versus Exponential Mechanism

- Let f be a function on tables that returns a real number.

- Define: score function w(D,O) = |f(D) - O|
- Sensitivity of w = $\max_{D,D'}$ (|f(D) – O| - |f(D') – O|)
  $\leq \max_{D,D'}$ |f(D) – f(D')| = sensitivity of f

- Exponential mechanisms returns an output f(D) + η with probability proportional to

$$e^{\frac{\varepsilon}{2\Delta} \cdot |f(D) - f(D) - \eta|}$$

Laplace noise with parameter 2Δ/ε

Summary of Exponential Mechanism

- Differential privacy for cases when output perturbation does not make sense.
- Idea: Make better outputs exponentially more likely; Sample from the resulting distribution.
- Every differentially private algorithm is captured by exponential mechanism.
  - By choosing the appropriate score function.

Summary of Exponential Mechanism

- Utility of the mechanism only depends on log(|Outputs|)
  - Can work well even if output space is exponential in the input

- However, sampling an output may not be computationally efficient if output space is large.

Duke UNIVERSITY

This class

- Exponential Mechanism: when the answer is not a real number

- Median Mechanism: Answering a stream of queries

Answering multiple queries

- Suppose total budget is ε.

- And each query uses δ privacy (in order to get utility)
  - Queries may be coming from different researchers
  - But they may collude …

- Then total number of queries answered is only k = ε/δ.

Answering correlated queries

- q1 = q2 = q3 = … = qk = “what fraction of the class is from China”?
- If we answer each query independently with Laplace mechanism, then we can’t answer any more queries.
- But, we could have just used Laplace mechanism once, and then reused the same answer for all the remaining queries.
  - We can still answer k-1 more queries!
- **Qn: can we figure out whether a query is “easy” – answerable from previous queries?**

Median Mechanism

- C_0 = set of all databases *// world consistent with existing query answers*
- Given a query q_i,
  - If q_i is a "hard" query:
    - Answer q_i using Laplace mechanism (a_i + noise)
    - Find S subset of C_i-1, such that for all D in S, $|f(D) – a_i| \leq \alpha/50$
    - $C_i = S$
  - If q_i is an "easy" query:
    - Compute $q_i(D)$ for all D in C_i-1
    - Return the median of all the computed $q_i(D)$
    - $C_i = C_{i-1}$

Median Mechanism

- When is a query "easy"?
  - When more than half the databases D' have |qi(D') – **qi(D)**| < ε
  - Then the median of all the answers is close to the true answer **ai = qi(D)**
  - But this could leak information …
  - Solution: Compute a noisy version of …

$$r_i = \frac{\sum_{S \in C_{i-1}} \exp(-\epsilon^{-1}|f_i(D) - f_i(S)|)}{|C_{i-1}|}.$$

Duke UNIVERSITY

Summary

- Exponential mechanism can be used to ensure differential privacy when range of algorithm is not a real number.
- Median mechanism can be used to answer streams of queries.

Duke
UNIVERSITY

Next class

- Smooth sensitivity and sampling