18734 Recitation

Basic Probability Theory

Laplace Mechanism

Definition of Probability

- ***Experiment***: toss a coin twice
- ***Sample space***: possible outcomes of an experiment
  - S = {HH, HT, TH, TT}
- ***Event***: a subset of possible outcomes
  - A={HH}, B={HT, TH}, C={TT}
- ***Probability of an event*** : an number assigned to an event Pr(A)
  - Axiom 1: Pr(A) ≥ 0
  - Axiom 2: Pr(S) = 1
  - Axiom 3: For every sequence of disjoint events

$$\Pr(\bigcup_i A_i) = \sum_i \Pr(A_i)$$

Slide: Rong Jin

Joint Probability

- For events A and B, **joint probability** $\Pr(A \cap B)$ stands for the probability that both events happen.
- Example: A={HH}, B={HT, TH}, what is the joint probability $\Pr(A \cap B)$?

Slide: Rong Jin

Conditional Probability

- If A and B are events with $\Pr(A) > 0$, the ***conditional probability of B given A*** is

$$\Pr(B|A) = \Pr(A \cap B)/\Pr(A)$$

Slide: Rong Jin

Random Variable (RV)

- A ***random variable X*** is a function from the sample space to a real number
- X : {HH} -> 2;   {HT, TH} -> 1;   {TT} -> 0

- Pr(X=0) = Pr(C), where C = {TT}

- A discrete RV takes on finite number of values
- A continuous RV can take an uncountable number of values

Discrete RV

- Probability Mass Function (PMF) p_X gives the probability that X will take on a particular value

- $p_X(x) = Pr(X=x)$

- $\sum_i p_X(x_i) = 1$

Continuous RV

- Probability Density Function (PDF) f_X is a non-negative function such that

$$Pr(a <= X <= b) = \int_a^b f_X(x)dx$$

- The integral from $-\infty$ to $+\infty$ is 1

- $Pr(X=a) = 0$

Probability Distribution

- A distribution assigns a probability to each event in the sample space

Discrete Uniform Distribution

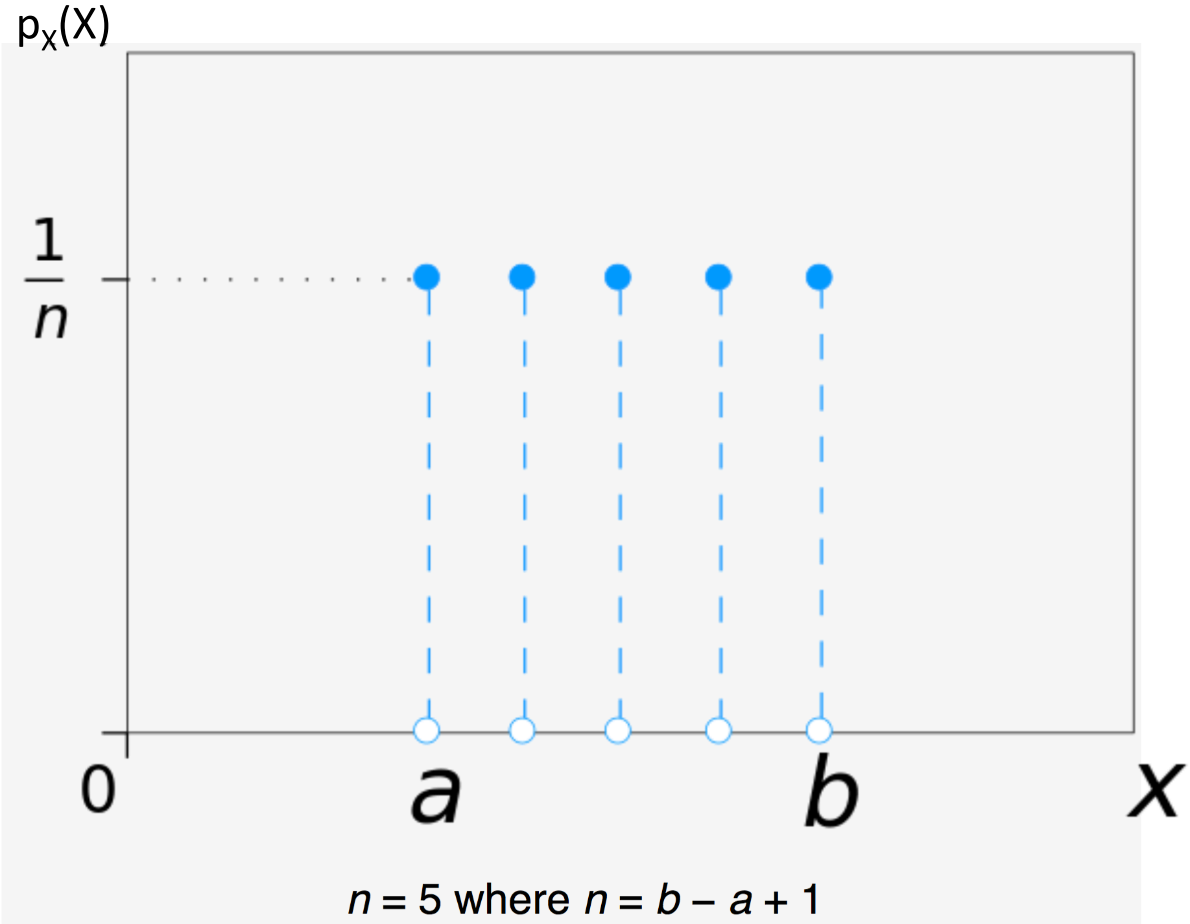

Continuous Uniform Distribution

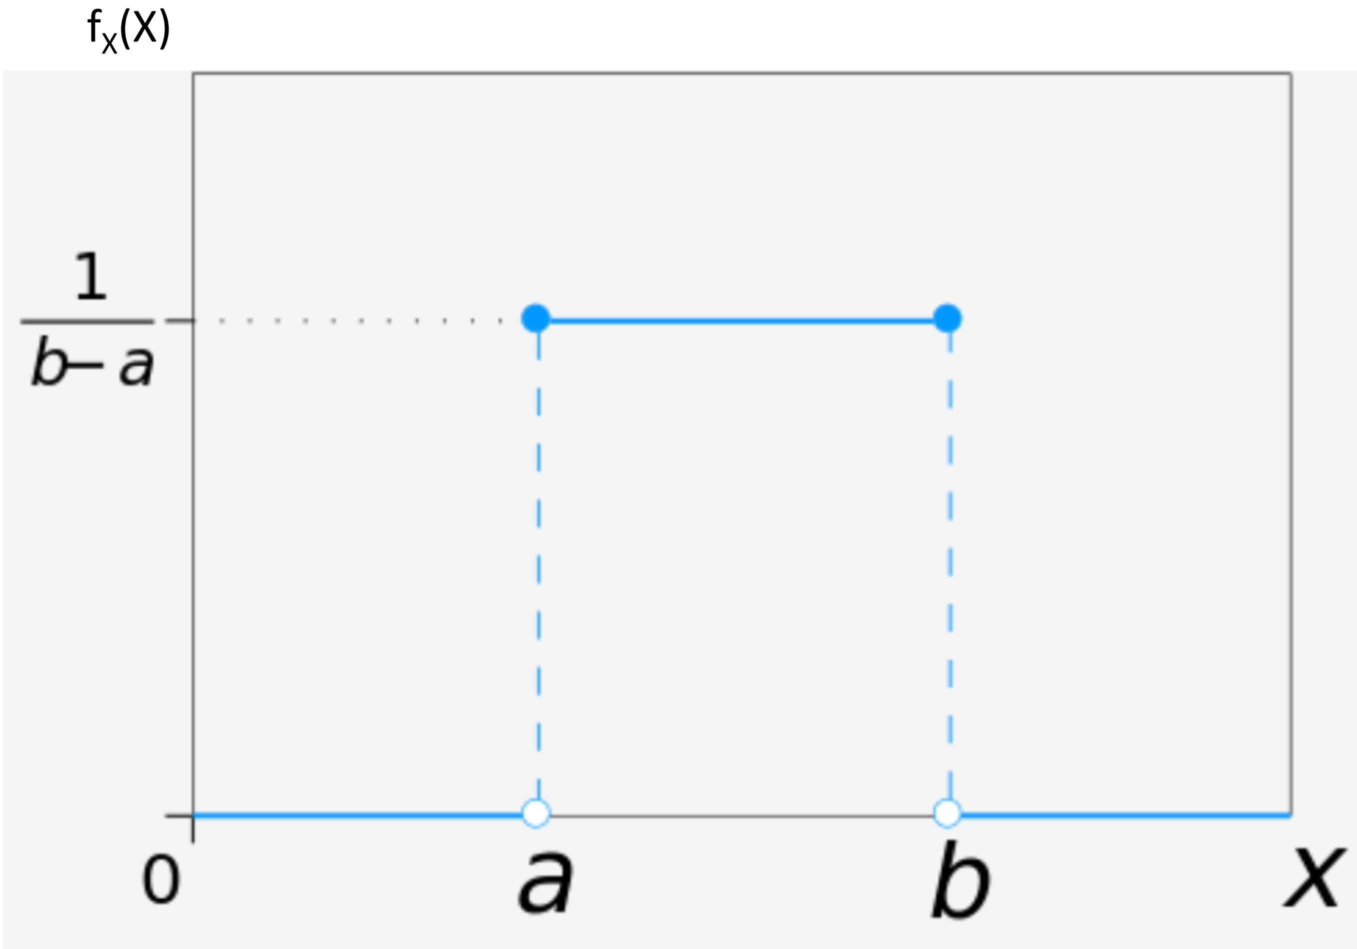

Laplace Distribution

PDF = $\frac{1}{2b} exp(-\frac{|y-\mu|}{b})$

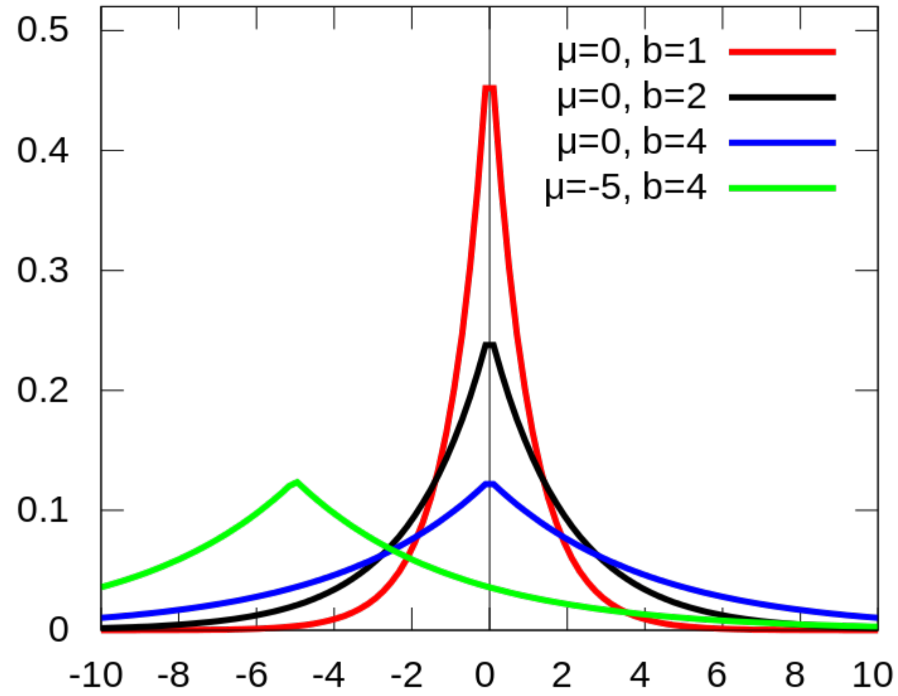

Source: Wikipedia

Laplace Distribution

- Laplace distribution $\mathrm{Lap}(\lambda)$ has density

$$h(y) \propto e^{-|y|/\lambda}$$

- Changing one point translates curve

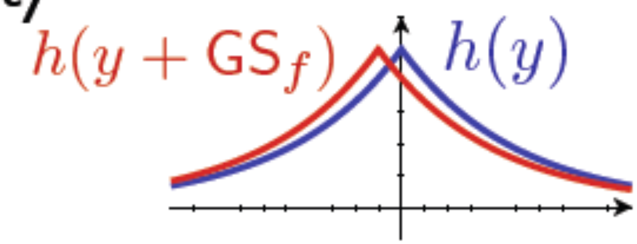

Change of notation from previous slide:

μ → 0

b → λ

Differential Privacy: Definition

Randomized function A has ε-differential privacy if for all data sets D_1 and D_2 differing by at most one element and all subsets S of the range of A,

$$\Pr[A(D_1) \in S] \leq e^{\varepsilon} \Pr[A(D_2) \in S]$$

Slide: Adam Smith

Laplace Mechanism

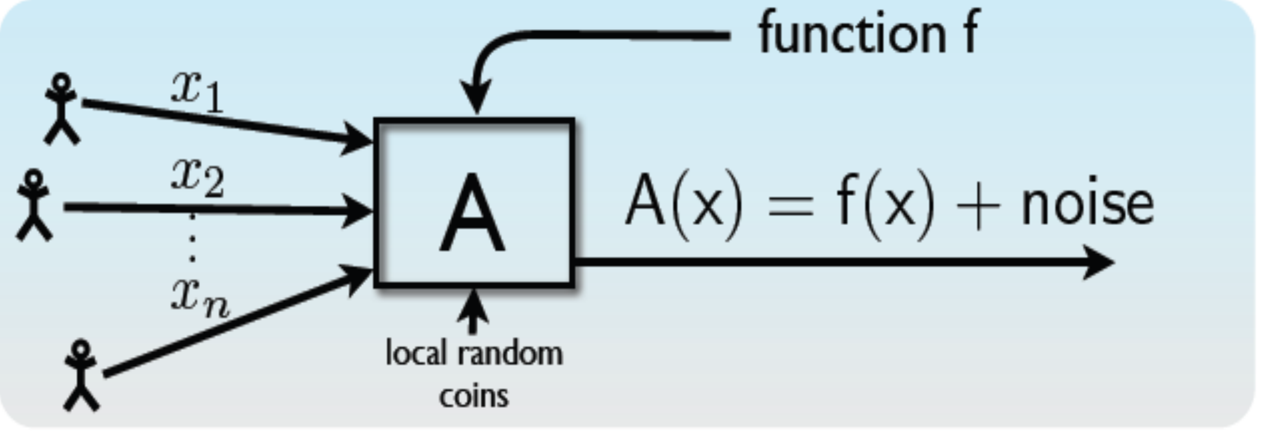

- Global Sensitivity: $\mathrm{GS}_f = \max_{\text{neighbors } x,x'} \|f(x) - f(x')\|_1$

Theorem: If $\mathrm{A}(\mathrm{x}) = f(\mathrm{x}) + \mathrm{Lap}\left(\frac{\mathrm{GS}_f}{\epsilon}\right)$, then A is ϵ-differentially private.

Laplace Mechanism: Proof Idea

Theorem: If $\mathsf{A}(\mathsf{x}) = f(\mathsf{x}) + \mathsf{Lap}\left(\frac{\mathsf{GS}_f}{\epsilon}\right)$, then A is ϵ-differentially private.

Laplace distribution $\mathsf{Lap}(\lambda)$ has density $h(y) \propto e^{-\frac{\|y\|_1}{\lambda}}$

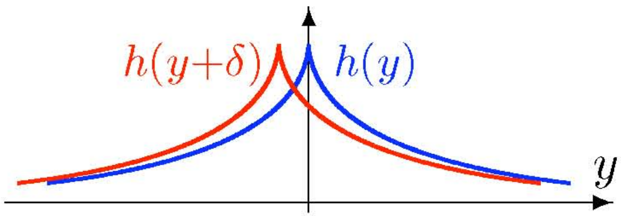

Sliding property of $\mathsf{Lap}\left(\frac{\mathsf{GS}_f}{\varepsilon}\right)$: $\frac{h(y)}{h(y+\delta)} \leq e^{\varepsilon \cdot \frac{\|\delta\|}{\mathsf{GS}_f}}$ for all y, δ

Proof idea:

$A(x)$: blue curve

$A(x')$: red curve

$\delta = f(x) - f(x') \leq \mathsf{GS}_f$

Laplace Mechanism: Proof

- To Prove:

$$\Pr[A(x) \in S\,] \leq e^{\varepsilon} \Pr[A(x') \in S\,]$$

- Distribution of A(x): Laplace(f(x), GS_f/λ)
- Distribution of A(x'): Laplace(f(x'), GS_f/λ)

Laplace Mechanism: Proof

$$\frac{Pr[A(x) \in S]}{Pr[A(x') \in S]}$$

$$= \frac{Pr[y \leq A(x) \leq y + dy]}{Pr[y \leq A(x') \leq y + dy]}$$

$$= \frac{e^{-|y-f(x)|/\lambda}}{e^{-|y-f(x')|/\lambda}}$$

$$= e^{\frac{|y-f(x')|-|y-f(x)|}{\lambda}}$$

$$\leq e^{\frac{|f(x)-f(x')|}{\lambda}} \leq e^{\frac{GS_f \epsilon}{GS_f}} = e^{\epsilon}$$