The Laplace Distribution

- randomly draw a value v with probability

$$\propto \frac{1}{2b} e^{-|v|/b}$$

($\propto$: proportional to)

- mean = max = 0
- larger $|v|$ less likely
- b is a parameter
- $b \to 0$: rapid fall-off in $|v|$
- $b \to \infty$: ~~slower fall-off~~

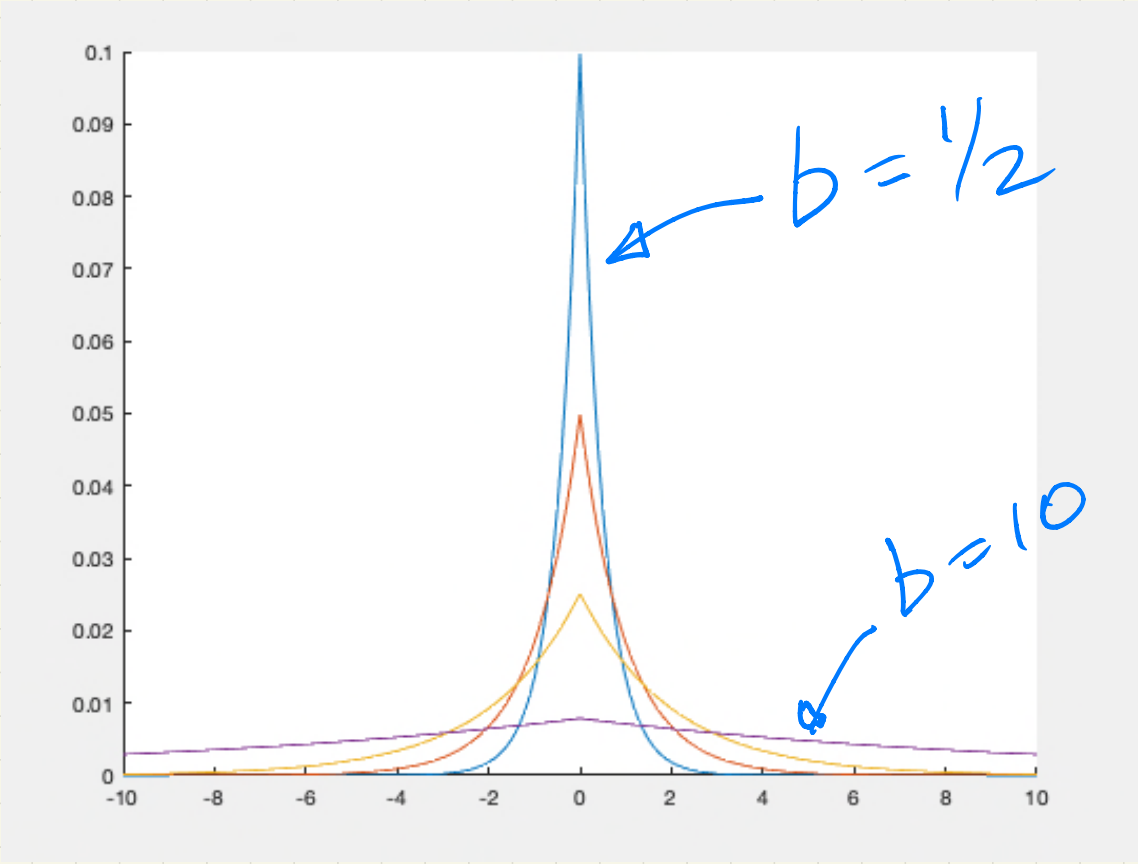
0.1
0.09
0.08
0.07
0.06
0.05
0.04
0.03
0.02
0.01
0
-10
-8
-6
-4
-2
0
2
4
6
8
10
b = 1/2
b = 10

Laplace Mechanism:

- compute $f(\bar{x})$ exactly
- output: $f(x) + v$

  where v is chosen randomly from Laplace distribution with $b = \Delta f / \varepsilon$
- smaller Δf/larger ε: less noise
- larger Δf/smaller ε: more noise

Privacy: Laplace Mech. satisfies ε-DP. (proof)

Utility:

$|\text{output} - f(\vec{x})| \approx$ stdev of $\nu \approx b = \frac{\Delta f}{\varepsilon}$

So if e.g. $\Delta f = 1/n$, we can set $\varepsilon = 1/\sqrt{n}$, then $b = \frac{1/n}{1/\sqrt{n}} = \frac{1}{\sqrt{n}}$.

If e.g. $\Delta f = 1/\sqrt{n}$, set $\varepsilon = 1/n^{1/4}$, then $b = \frac{1/n^{1/2}}{1/n^{1/4}} = 1/n^{1/4}$

In general, if Δf diminishes with n, can choose ε s.t.

① $\varepsilon \to 0$ as $n \to \infty$

perfect privacy

② $|\text{output} - f(\vec{x})| \to 0$ as $n \to \infty$

perfect accuracy

Proof of Privacy:

- fix neighboring $\vec{x}, \vec{x}'$
- let A denote Laplace Mech.
- let y be any possibe output

Then:

$$\frac{Pr[A(\vec{x})=y]}{Pr[A(\vec{x}')=y]} = \frac{\frac{1}{2b}e^{-|f(\vec{x})-y|/b}}{\frac{1}{2b}e^{-|f(\vec{x}')-y|/b}}$$

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

call this z. claim:

$$z \leq |f(\vec{x}')-f(\vec{x})|$$

Case 1:

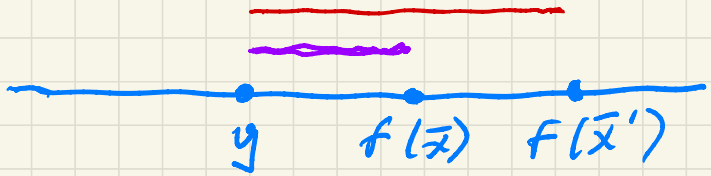

— minus — $= |f(x') - f(x)|$ ✓

Case 2:

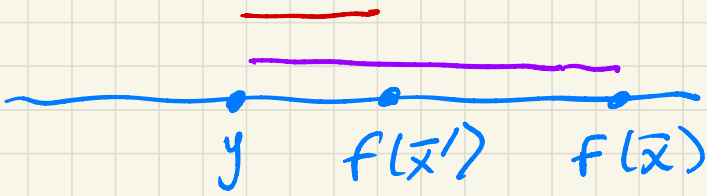

— minus — $= -|f(x') - f(x)|$ ✓

Case 3:

— minus — $\leq |f(x') - f(x)|$ ✓

So then

$$e^{(|f(x')-y|-|f(x)-y|)/b}$$

$$\leq e^{|f(x')-f(x)|/b}$$

$$\leq e^{\Delta f/b} \quad \text{(sensitivity)}$$

$$= e^{\Delta f/(\Delta f/\varepsilon)} \quad (b = \Delta f/\varepsilon)$$

$$= e^{\varepsilon}.$$

Now if S is any set of outputs $S = \{y_1, y_2, \dots\}$

$$\Pr[A(\bar{x}) \in S] = \sum_{y \in S} \Pr[A(\bar{x}) = y]$$

$$\leq \sum_{y \in S} e^{\varepsilon} \Pr[A(\bar{x}') = y]$$

$$= e^{\varepsilon} \sum_{y \in S} \Pr[A(\bar{x}') = y]$$

$$= e^{\varepsilon} \Pr[A(\bar{x}') \in S].$$

More General Methods?

- Laplace great for computing functions with numeric inputs/outputs
- What about things like:

Input = medical DB, output = neural network

Input = social network, output = clustering of users

Input = location & health data, output = COVID-19 "hot spots"

A more general framework:

- Inputs x (med DBs, social NW, etc)
- outputs or "solutions" y (NNs, clustering, etc.)

Then for <u>any</u> pair (x,y),

let $u(x,y)$ be a function expressing how "good" output y is for input x.

u for "utility"

Examples

- u(med DB D, neural net N) =
  error of N on D in predicting disease
- u(soc. net. S, clustering C) =
  measure of how well C divides S into similar groups
- u(loc/health data D, set of locations S) =
  measure of how well S captures density of infected individuals

Generalized Sensitivity

Define sensitivity of $u(x,y)$:

$$\Delta u = \max_{\text{outputs } y} \max_{\text{neighboring inputs } x,x'} |u(x,y) - u(x',y)|$$

How much can the utility of y change by removing your data from input?

A More General DP Method:

The Exponential Mechanism

"Algorithm": On input x, pick output y randomly according to the distribution:

$$\Pr[y|x] \propto e^{\varepsilon \cdot u(x,y)/2\Delta u}$$

- input x is given
- utility function is chosen
  - → determines Δu
- ε is privacy parameter

Properties of Exp. Mech.

Privacy: Exp. Mech is ε-DP.

(proof: similar to Laplace)

Utility: For input x, let

$$y_x^* = \arg\max_y \{u(x,y)\}$$

i.e. y_x^* is best output/soln for x.

Then w.h.p. Exp. Mech. outputs y s.t.

$$u(x,y) \geq u(x,y_x^*) - O\left(\frac{\Delta u}{\varepsilon}\right).$$

Remarks

- As with Laplace, if Δu is a diminishing function of n, we can approach perfect privacy and perfect utility as $n \to \infty$.
- "Algorithm": may be difficult or impossible to sample distribution!
- Proof of concept, need bespoke DP algo design