CS208: Applied Privacy for Data Science
DP Foundations: the Laplace Mechanism

James Honaker & Salil Vadhan

School of Engineering & Applied Sciences
Harvard University

February 22, 2019

Center for Research on Computation and Society
at Harvard John A. Paulson School of Engineering and Applied Sciences

Differential Privacy

M is ϵ-DP if

$$Pr[M(D,q) \in T] \leq (1+\epsilon) Pr[M(D',q) \in T], \qquad \forall\, T, q.$$

Differential Privacy

M is ϵ-DP if

$$Pr[M(D,q) \in T] \leq (1+\epsilon)Pr[M(D',q) \in T], \qquad \forall\, T, q.$$

- D, D' Neighbouring datasets

Differential Privacy

M is ϵ-DP if

$$Pr[M(D,q) \in T] \leq e^{\epsilon} Pr[M(D',q) \in T], \qquad \forall\ T, q.$$

- D, D' Neighbouring datasets
- M Mechanism that Maps from data to result

Differential Privacy

M is ϵ-DP if

$$Pr[M(D,q) \in T] \leq e^{\epsilon} Pr[M(D',q) \in T], \qquad \forall\, T, q.$$

- D, D' Neighbouring datasets
- M Mechanism that Maps from data to result
- q Query
- T Set providing a decision rule

e^{ϵ} vs. $(1+\epsilon)$

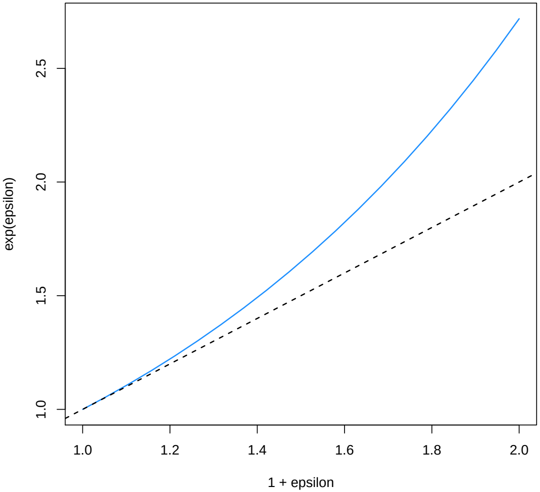

see `expEpsilon.r`

$x_i = x_{min}$

$\bar{X}(x_i) = \sum_{j \neq i} \frac{x_j}{N-1} + \frac{x_{min}}{N}$

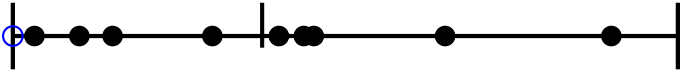

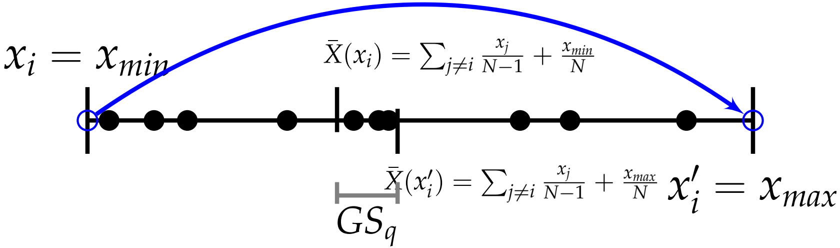
$x_i = x_{min}$
$\bar{X}(x_i) = \sum_{j\neq i} \frac{x_j}{N-1} + \frac{x_{min}}{N}$
$\bar{X}(x_i') = \sum_{j\neq i} \frac{x_j}{N-1} + \frac{x_{max}}{N}$
$x_i' = x_{max}$
GS_q

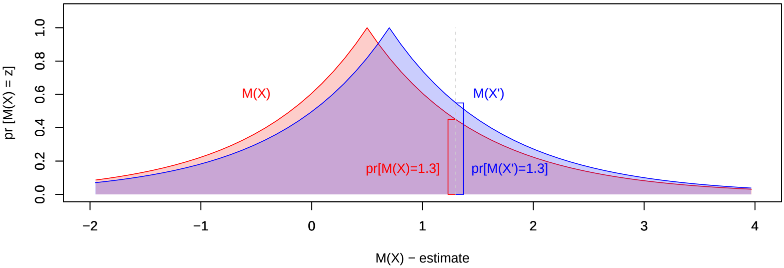

M(X)
M(X')
pr[M(X)=1.3]
pr[M(X')=1.3]
pr [M(X) = z]
M(X) − estimate
0.0
0.2
0.4
0.6
0.8
1.0
−2
−1
0
1
2
3
4

The Laplace has density over y:

$$f_{Laplace}(y|s,\mu) = \text{Lap}(s,\mu) = \frac{1}{2s}\exp\Big(-\frac{|y-\mu|}{s}\Big)$$

The Laplace has density over y:

$$f_{Laplace}(y|s) = \text{Lap}(s) = \frac{1}{2s} \exp\Big(-\frac{|y|}{s} \Big)$$

We were given the theorem:

$$M(x, q) = q(x) + \text{Lap}(GS_q/\epsilon)$$

The Laplace has density over y:

$$f_{Laplace}(y|s) = \text{Lap}(s) = \frac{1}{2s}\exp\Big(-\frac{|y|}{s}\Big)$$

We were given the theorem:

$$M(x, q) = q(x) + \text{Lap}(GS_q/\epsilon)$$

So our differentially private mean, $M(X)$, which combines the "true" sample mean with Laplace noise, becomes:

$$M(x) = \bar{x} + Z; \qquad Z \sim \text{Lap}(s = GS_q/\epsilon)$$

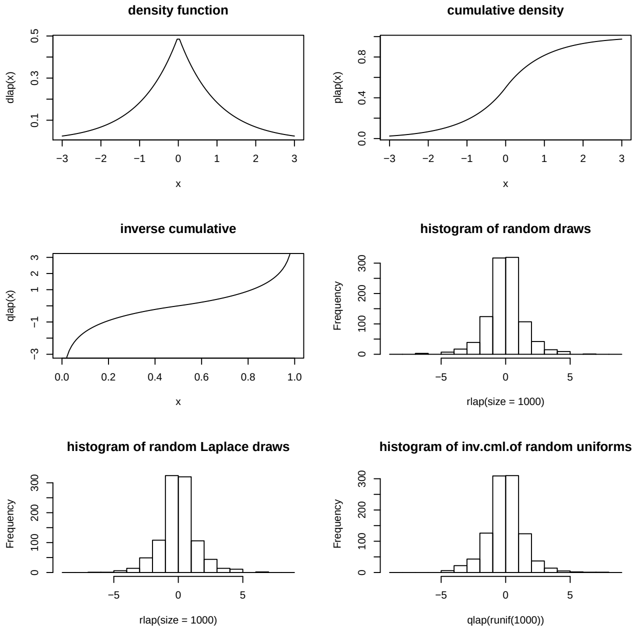

see `laplaceDistributions.r`

$$\frac{pr[M(x)=t]}{pr[M(x')=t]} = \frac{e^{\frac{-\epsilon|\bar{x}-t|}{GS_q}}}{e^{\frac{-\epsilon|\bar{x}'-t|}{GS_q}}} = e^{\frac{\epsilon|\bar{x}'-t|-\epsilon|\bar{X}-t|}{GS_q}} = e^{\frac{\epsilon|\bar{x}'-\bar{x}|}{GS_q}} \leq e^{\epsilon}$$

since we know $GS_q \geq |\bar{x}' - \bar{x}|$ by the def. of sensitivity. Thus we meet the original definition:

$$Pr[M(x)=t] \leq e^{\epsilon} Pr[M(x')=t]$$

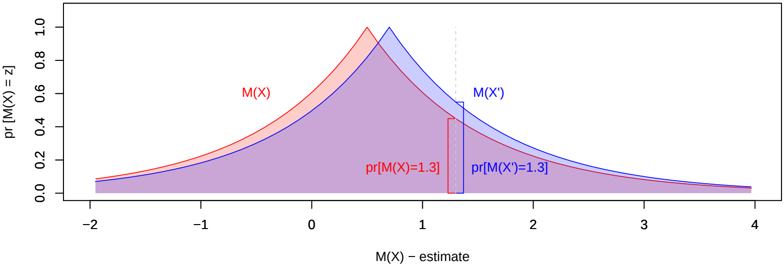

Two Laplace distributions, for two adjacent datasets x and x'. The definition of ϵ-differential privacy requires the ratio of $M(x)/M(x')$ is not greater than e^{ϵ} for all points along the x-axis. Thus for any realized output z (for example here, $z = 1.3$) we can not determine that x or x' were more likely to have produced z.

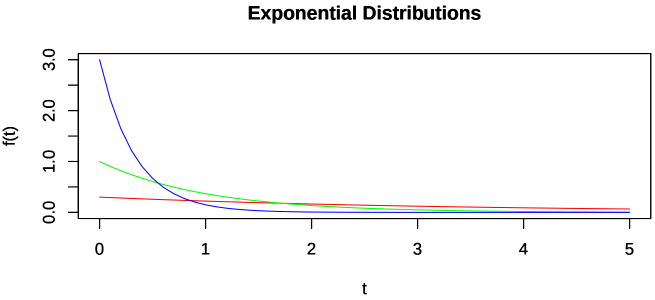
Exponential Distributions
f(t)
0.0
1.0
2.0
3.0
0
1
2
3
4
5
t

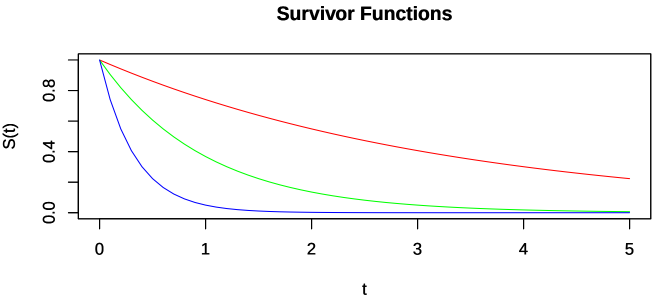
Survivor Functions
S(t)
0.0
0.4
0.8
0
1
2
3
4
5
t

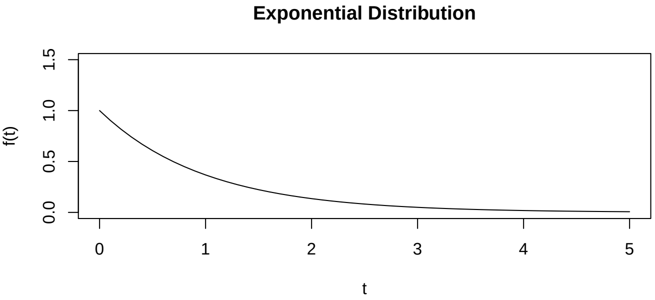
Exponential Distribution
f(t)
0.0
0.5
1.0
1.5
0
1
2
3
4
5
t

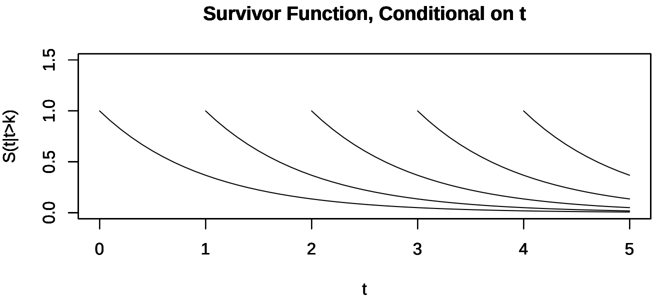
Survivor Function, Conditional on t
S(t|t>k)
0.0
0.5
1.0
1.5
0
1
2
3
4
5
t

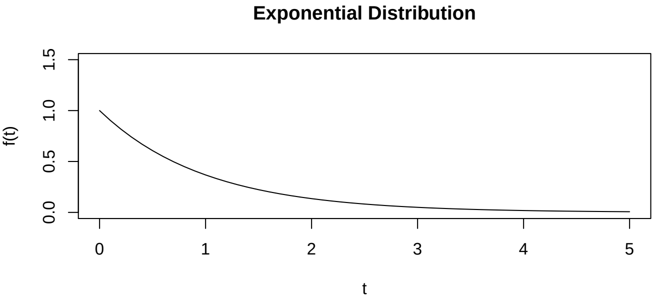
Exponential Distribution
f(t)
0.0
0.5
1.0
1.5
0
1
2
3
4
5
t

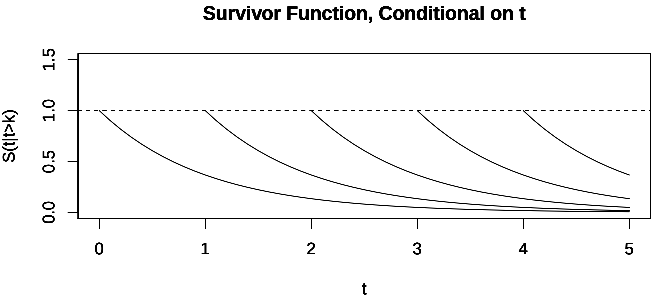
Survivor Function, Conditional on t
S(t|t>k)
0.0
0.5
1.0
1.5
0
1
2
3
4
5
t

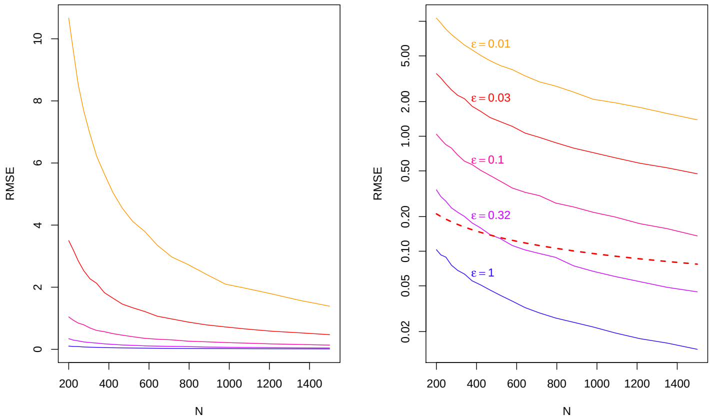

see `laplaceMeanRelease.r`