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

School of Engineering & Applied Sciences
Harvard University

February 24, 2025

Wikimedia

page_id	project	country	date	count
23110294	en.wikipedia.org	CH	2022-08-15	42
28278	fr.wikipedia.org	US	2022-08-15	17
...	...	...	...	...

Discussion

Assume the DP system for releasing DP page/country counts works as described. Should a user who visits a page have the ability to opt out of having their data used?

- Can data be the price of admission?
- If DP states the answer doesn't change if a user opts out, why not allow a user to opt out?
- Can't sophisticated users opt out anyway? Is that fair?

Wikimedia

page_id	project	country	date	count
23110294	en.wikipedia.org	CH	2022-08-15	42
28278	fr.wikipedia.org	US	2022-08-15	17
...	...	...	...	...

Wikimedia

- “**contribution bounding:** [DP] must hide all the contributions of a single Wikipedia user during a single day…we must truncate the input data to ensure that each user does not contribute more than a certain amount of data. ”
- “the data is grouped by page and country, using **the public data** to generate the list of possible counts. Listing these groups is an important step to achieve differential privacy

Wikimedia

- "**contribution bounding:** [DP] must hide all the contributions of a single Wikipedia user during a single day… we must truncate the input data to ensure that each user does not contribute more than a certain amount of data. "
- "the data is grouped by page and country, using **the public data** to generate the list of possible counts. Listing these groups is an important step to achieve differential privacy: if we simply use the counts that appear in the data, this can break privacy guarantees."
- "**spurious data:** for these [zero] counts, noise was added to a count of 0. To prevent too many of these to end up in the real data, we add a post-processing step to remove the counts below a certain threshold."

Approximate DP

Approximate Differential Privacy

Def: M is (ε, δ)-DP if for all $x \sim x'$, we have

$$\forall T \quad \Pr[M(x) \in T] \leq \mathrm{e}^{\varepsilon} \cdot \Pr[M(x') \in T] + \delta$$

Approximate DP

Benefits of Approximate DP

- More mechanisms, e.g. Gaussian Mechanism:

$$M(x,q) = q(x) + \mathcal{N}\left(0, \sigma^2\right),$$

$$\text{for } \sigma = \frac{\mathrm{GS}_q}{\varepsilon} \cdot \sqrt{2\ln(2/\delta)}$$

zero-Concentrated DP (zCDP)

ρ-zCDP: privacy loss is “subGaussian” – dominated by a Gaussian r.v. with mean ρ and variance 2ρ

- ε-DP implies $(\varepsilon^2/2)$-zCDP
- ρ-zCDP implies $\left(\rho + 2\sqrt{\rho \log(1/\delta)}, \delta\right)$-DP for all δ
- Composition: $\rho_i's$ add up
- Gaussian mechanism:
  $M(x, q) = q(x) + \mathcal{N}(0, (\Delta q)^2/2\rho\,)$ is ρ-zCDP

$$\rho = 0.15; \delta = 10^{-7} \Rightarrow .15 + 2\sqrt{.15\, log(1/10^{-7})}$$

Gaussian Mechanism

$$M(x, q) = q(x) + \mathcal{N}(0, \sigma^2),$$

$$\text{for} \quad \sigma = \frac{GS_q}{\epsilon}\sqrt{2\ln(2/\delta)}.$$

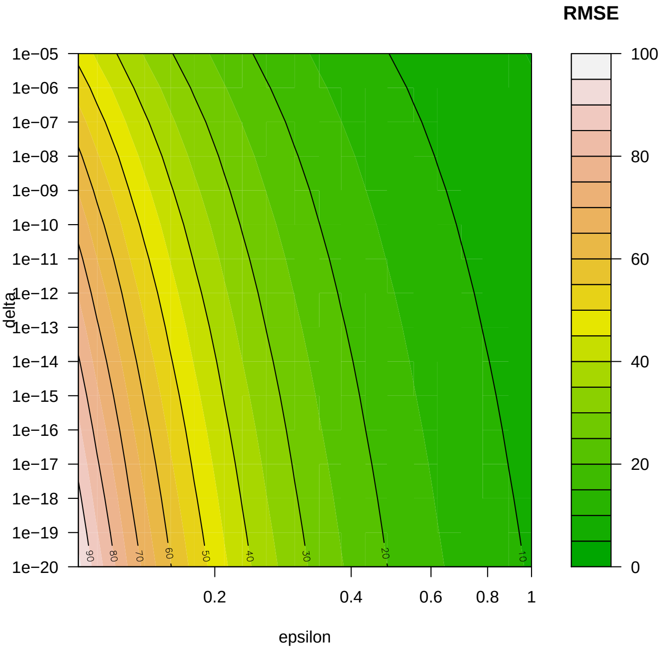

RMSE
100
80
60
40
20
0
1e−05
1e−06
1e−07
1e−08
1e−09
1e−10
1e−11
1e−12
1e−13
1e−14
1e−15
1e−16
1e−17
1e−18
1e−19
1e−20
delta
0.2
0.4
0.6
0.8
1
epsilon
90
80
70
60
50
40
30
20
10

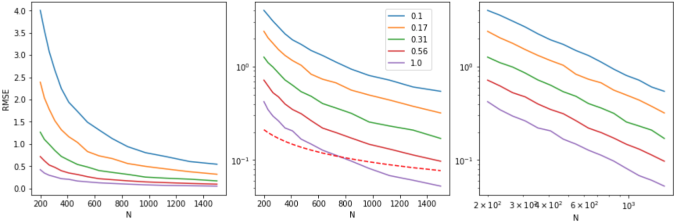

Gaussian Mechanism

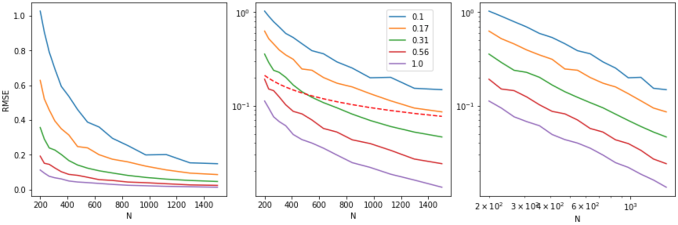

Laplace Mechanism

Properties of the Definition

- Suffices to check pointwise: M is ϵ-DP if and only if
$$\forall x \sim x', \forall q, \forall t \ \Pr[M(x,q) = t] \leq e^{\epsilon} \cdot \Pr[M(x',q) = t]$$
Replace with densities for continuous distributions

- Closed under post-processing: if M is ϵ-DP and f is any function, then $M'(x,q) = f(M(x,q))$ is also ϵ-DP.

- (Basic) composition: If M_i is ϵ_i-DP for $i = 1, \ldots, k$, then
$$M\big(x, (q_1, \ldots, q_k)\big) = \big(M_1(x, q_1), \ldots, M_k(x, q_k)\big)$$
is $(\epsilon_1 + \cdots + \epsilon_k)$-DP.
  - Use independent randomness for k queries.
  - Holds even if q_i's are adaptively chosen by an adversary.

Confidence Interval Construction

Given an estimate $\hat{y}$, of a quantity y^*, a confidence interval, $\mathrm{ci}(y^*|\hat{y}, \alpha) = [ci_{lower}, ci_{upper}]$ often simply $\mathrm{ci}_{1-\alpha}(y^*)$, has *proper coverage* if:

$$\mathrm{Prob}[y^* \in [ci_{lower}, ci_{upper}]] = 1 - \alpha$$

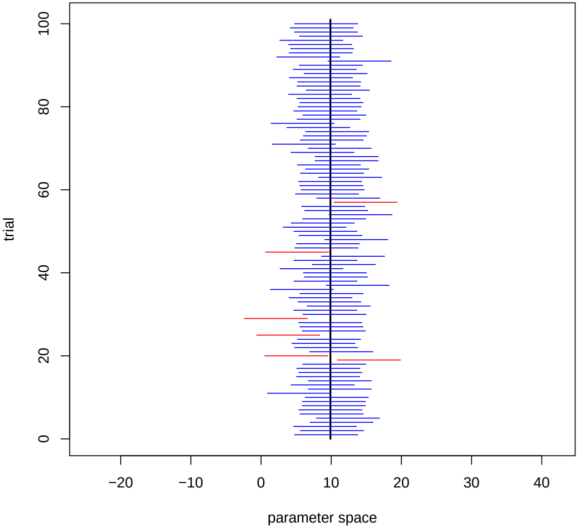
trial
100
80
60
40
20
0
−20
−10
0
10
20
30
40
parameter space

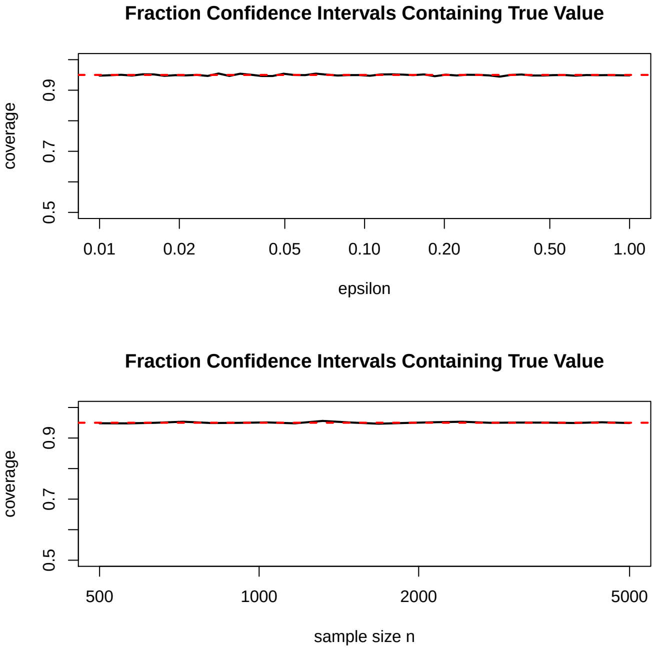

Fraction Confidence Intervals Containing True Value
coverage
0.5
0.7
0.9
0.01
0.02
0.05
0.10
0.20
0.50
1.00
epsilon
Fraction Confidence Intervals Containing True Value
coverage
0.5
0.7
0.9
500
1000
2000
5000
sample size n

Education Values

Codebook for Census PUMS 5 Percent CS208 Datasets

`educ`	1:	No schooling completed,
	2:	Nursery school to 4th grade,
	3:	5th grade or 6th grade,
	4:	7th grade or 8th grade,
	5:	9th grade,
	6:	10th grade,
	7:	11th grade,
	8:	12th grade, no diploma,
	9:	High school graduate,
	10:	Some college, but less than 1 year,
	11:	One or more years of college, no degree,
	12:	Associate degree,
	13:	Bachelor's degree,
	14:	Master's degree,
	15:	Professional degree,
	16:	Doctorate degree.

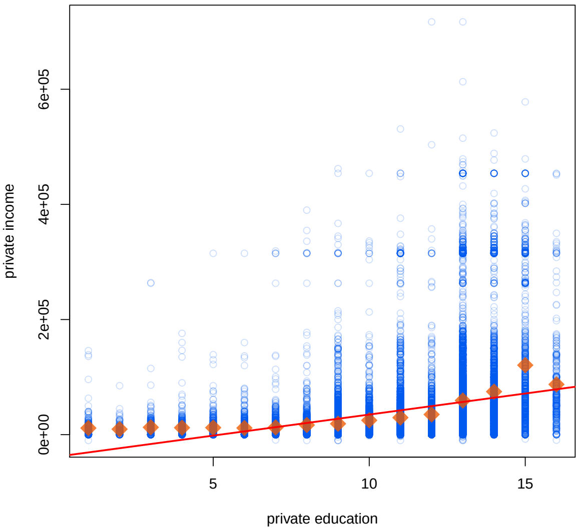

Positive relationship of education and income in PUMS

What is the global sensitivity of the sample correlation?

$$\mathrm{corr}(X, Y) = \frac{1}{N-1} \sum_{i=1}^{N} \frac{(x_i - \bar{x})(y_i - \bar{y})}{s_x s_y}$$