Less is More: Revisiting the Gaussian Mechanism for Differential Privacy

Tianxi Ji[†] and Pan Li[*]

[†]Texas Tech University
[*]Case Western Reserve University

33rd USENIX Security Symposium
Philadelphia, PA, USA
August 14-16, 2024

Contributions in a nutshell

- Geometric representation of the privacy loss in DP
- Approaches to design new DP mechanisms by leveraging measure concentration of random triangles

Outline

- Background
- The curse of full-rank covariance matrices
- A new mechanism (R1SMG)
- Open problems

Differential Privacy (DP) Background

(ϵ, δ)-DP is recognized as the fundamental building block for privacy-preserving database query, data mining, learning...

Definition

A randomized mechanism $\mathcal{M}$ satisfies (ϵ, δ)-DP if for any two neighboring datasets, $\boldsymbol{x}, \boldsymbol{x}'$, and any outcome $s \in \mathcal{S} \subseteq \mathrm{Range}(\mathcal{M})$,

$$\Pr[\mathrm{PLRV} \geq \epsilon] \leq \delta$$

holds, where $\mathrm{PLRV} = \ln\left(\frac{\Pr[\mathcal{M}(\boldsymbol{x})=s]}{\Pr[\mathcal{M}(\boldsymbol{x}')=s]}\right)$, $\epsilon > 0$ and $0 < \delta \ll 1$.

Differential Privacy (DP) Background

The Gaussian mechanism is an essential tool to achieve (ϵ, δ)-DP for a given computation $f(\mathbf{x}) \in \mathbb{R}^{M \times N}$, $M \geq 1, N \geq 1$.

Variants of the mechanism:

- The classic Gaussian mechanism adds $\mathcal{N}(0, \sigma_C^2)$ to $f(\mathbf{x})$ [Dwork et al., EUROCRYPT 2006]
- The analytic Gaussian mechanism adds $\mathcal{N}(0, \sigma_A^2)$ to $f(\mathbf{x})$ [Balle and Wang, ICML 2018]
- The MVG mechanism adds $\mathcal{N}_{M,N}(\mathbf{0}, \boldsymbol{\Sigma}, \boldsymbol{\Psi})$ to $f(\mathbf{x})$ [Chanyaswa et al., CCS 2018]

σ_C^2, σ_A^2, $\boldsymbol{\Sigma}$, and $\boldsymbol{\Psi}$ are all calibrated by ϵ, δ, and sensitivity $\Delta_2 f = \max_{\boldsymbol{x} \sim \boldsymbol{x}'} ||f(\boldsymbol{x}) - f(\boldsymbol{x}')||_2$

Outline

- Background
- The curse of full-rank covariance matrices
- A new mechanism (R1SMG)
- Open problems

The curse of full-rank noise covariance matrices

Define accuracy loss for mechanism $\mathcal{M}$

$$\mathcal{L} = ||\mathcal{M}(f(\boldsymbol{x})) - f(\boldsymbol{x})||_2^2 = ||\mathbf{n}||_2^2.$$

When $\mathcal{M}$ is the classic Gaussian mechanism (or its variants)

$$\mathbb{E}[\mathcal{L}] = Tr[Cov(\mathbf{n})].$$

Theorem

$$f(\mathbf{x}) \in \mathbb{R}^M, \mathbb{E}_{classic}[\mathcal{L}] = Tr[\sigma^2 \mathbf{I}_{M\times M}] \geq C_C(\Delta_2 f)^2, C_C = \frac{2\ln\left(\frac{1.25}{\delta}\right)}{\epsilon^2} M$$

$$f(\mathbf{x}) \in \mathbb{R}^M, \mathbb{E}_{analytic}[\mathcal{L}] = Tr[\sigma_A^2 \mathbf{I}_{M\times M}] \geq C_A(\Delta_2 f)^2, C_A = \frac{\left(\Phi^{-1}(\delta)\right)^2 + \epsilon}{\epsilon^2} M$$

$$f(\mathbf{x}) \in \mathbb{R}^{M\times N}, \mathbb{E}_{MVG}[\mathcal{L}] = Tr[\boldsymbol{\Sigma} \otimes \boldsymbol{\Psi}] \geq C_M(\Delta_2 f)^2, C_M = \frac{(\frac{5}{4}H_r + \frac{1}{4}H_{r,\frac{1}{2}})}{2\epsilon} MN$$

Curse: $\mathbb{E}[\mathcal{L}]$ is on the order of the dimension of $f(\mathbf{x})$

A hidden clue to lift the curse

"The algorithmic foundations of differential privacy" by Dwork and Roth (p. 261-265)

- $f(\cdot)$ is a query function, i.e., $f : \boldsymbol{x} \in \mathbb{N}^{|\mathcal{X}|} \rightarrow \mathbb{R}^M$
- Interested in $\mathcal{N}(0, \sigma^2 \mathbf{I})$ that can obscure $\boldsymbol{v} \triangleq f(\boldsymbol{x}) - f(\boldsymbol{x}')$
- $\mathcal{N}(0, \sigma^2 \mathbf{I})$ is spherically symmetric; represent the noise $\mathbf{n}$ using any fixed orthonormal basis $\boldsymbol{b}_1, \boldsymbol{b}_2, \cdots, \boldsymbol{b}_m$, i.e, $\mathbf{n} = \sum_{i=1}^{M} \lambda_i \boldsymbol{b}_i$, $\lambda_i \sim \mathcal{N}(0, \sigma^2), i \in [1, M]$
- WLOG, assume $\boldsymbol{b}_1$ is parallel to $\boldsymbol{v}$. Consequently,

$$\mathrm{PLRV}^{(\mathbf{s})}_{(\mathcal{G}(\boldsymbol{x})||\mathcal{G}(\boldsymbol{x}'))} = \left| \frac{1}{2\sigma^2} \left(||\boldsymbol{n}||^2 - ||\boldsymbol{n} + \boldsymbol{v}||^2 \right) \right| \ldots \leq \frac{1}{2\sigma^2} \left((\Delta_2 f)^2 + 2\lambda_1 \Delta_2 f \right)$$

A hidden Clue: PLRV is only related to λ_1 and $\Delta_2 f$

Outline

- Background
- The curse of full-rank covariance matrices
- A new mechanism (R1SMG)
- Open problems

The R1SMG mechanism—Design

Recall the clue: multivariate Gaussian noise whose covariance matrix has rank-1 is sufficient to achieve (ϵ, δ)-DP

The R1SMG Mechanism

For an arbitrary M-dimensional query function, $f(\boldsymbol{x}) \in \mathbb{R}^M$, the R1SMG mechanism is defined as

$$\mathcal{M}_{R1SMG}(f(\boldsymbol{x})) = f(\boldsymbol{x}) + \mathbf{n}$$
$$\mathbf{n} = \mathbf{v}\sqrt{\sigma_*}z, \quad \text{where} \quad z \sim \mathcal{N}(0,1), \quad \mathbf{v} \sim \mathbb{S}^{M-1}$$

$\mathbf{v}$ uniformly sampled from the unite sphere $\mathbb{S}^{M-1}$ embedded in $\mathbb{R}^M$.

$\mathbf{v}$ is random

Make PLRV well-defined.
Prevent privacy leakage of utilizing vector in the null space of $\mathbf{v}$.

The R1SMG mechanism—Privacy Guarantee

Theorem

The R1SMG mechanism achieves (ϵ, δ)-DP when $M > 2$, if $\sigma_* \geq \frac{2(\Delta_2 f)^2}{\epsilon\psi}$ where $\psi = \left(\frac{\delta\Gamma(\frac{M-1}{2})}{\sqrt{\pi}\Gamma(\frac{M}{2})}\right)^{\frac{2}{M-2}}$, and $\Gamma(\cdot)$ is the Gamma function.

Proof sketch: use (i) measure concentration of the random angle formed by random noise vectors **n** and **n**′ and (ii) law of sine

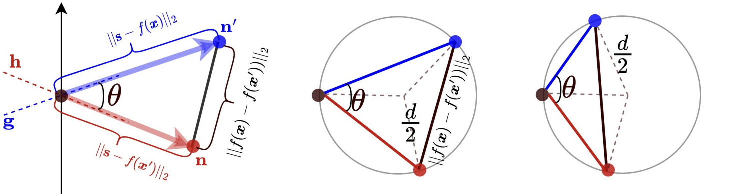

(a) geometric interpretation

(b) example of circumcircle with $\theta < \frac{\pi}{2}$

(c) example of circumcircle with $\theta > \frac{\pi}{2}$

$$\text{PLRV} \leq \frac{2}{\sigma_*}\left(\frac{\Delta_2 f}{\sin(\theta)}\right)^2$$

The R1SMG mechanism—Expected accuracy loss

Measure concentration: $\Pr\left[\left|\theta - \frac{\pi}{2}\right| \geq \theta_0\right] \leq \cdots = \delta$

θ converges to $\frac{\pi}{2}$ when dimension approaches infinity

Theorem (Less is more. Hide in the crowd.)

For any fixed feasible $\epsilon > 0, 0 < \delta < 1$, given a query result $f(\boldsymbol{x}) \in \mathbb{R}^M$, $\mathbb{E}_{R1SMG}[\mathcal{L}]$ has a decreasing trend as M increases. When M approaches infinity, $\mathbb{E}_{R1SMG}[\mathcal{L}]$ can be as low as $\frac{2(\Delta f)^2}{\epsilon}$.

Accuracy loss for a mechanism: $\mathcal{L} = ||\mathcal{M}(f(\boldsymbol{x})) - f(\boldsymbol{x})||_2^2 = ||\mathbf{n}||_2^2$.

$\mathcal{L}$ with both **larger kurtosis and skewness is preferred**

- Kurtosis, a descriptor of "tail extremity" of a probability distribution, defined as $\frac{\mathbb{E}[\mathcal{L}^4]}{(\mathbb{E}[\mathcal{L}^2])^2}$. A larger kurtosis means that extreme large values are less likely to be generated
- Skewness, a descriptor of the "bulk" of a probability distribution, defined as $\frac{\mathbb{E}[\mathcal{L}^3]}{(\mathbb{E}[\mathcal{L}^2])^{3/2}}$. A larger skewness means that the bulk of the samples is at the left region of the PDF

Theorem

The kurtosis and skewness of $\mathcal{L}$ in R1SMG is the largest.

The R1SMG mechanism—Caveat

Recall: the classic Gaussian mechanism requires $\epsilon < 1$ to obscure an arbitrary $\boldsymbol{v} = f(\boldsymbol{x}) - f(\boldsymbol{x}')$

Geometric interpretation:

- if $\epsilon < 1$, the noise components along the direction of $\boldsymbol{v}$ are also sufficient to obscure the difference, i.e., $\lambda_1 \boldsymbol{b}_1 - (\boldsymbol{b}_1^T \boldsymbol{n}')\boldsymbol{b}_1 = \boldsymbol{v}$
- if ϵ exceeds the upper bound, the magnitude of $\boldsymbol{n}$ and $\boldsymbol{n}'$ might be too small to obscure $\boldsymbol{v}$, since $||\boldsymbol{n}||, ||\boldsymbol{n}'|| \propto \frac{1}{\epsilon}$

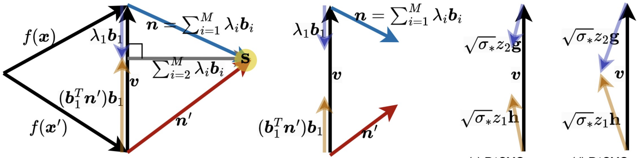

(a) classic Gaussian, $\epsilon < 1$ (b) classic Gaussian, ϵ too large (c) R1SMG, ϵ too large (d) R1SMG, ϵ small

Rule of thumb: $\epsilon < \frac{1}{M}\epsilon_{classic}$ (the exact bound is on our to-do list)

The R1SMG mechanism—Case study

Release the counts of Uber pickups in NYC from "4/1/2014 00:11:00" to "4/3/2014 23:57:00" in a DP manner. $f(\boldsymbol{x}) \in \mathbb{R}^{89\times 89}$

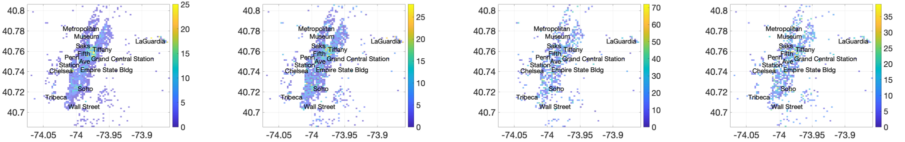

Original | R1SMG, $\epsilon = 10^{-5}$ | classic, $\epsilon = 0.5$ | Analytic, $\epsilon = 0.5$

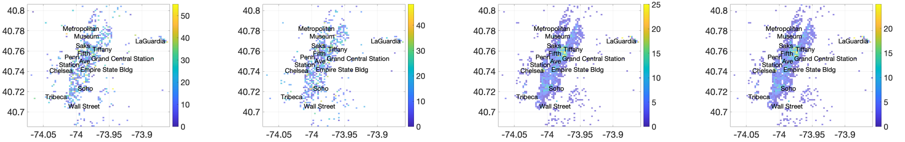

MVG, $\epsilon = 0.5$ | MGM, $\epsilon = 0.5$ | DAWA, $\epsilon = 0.5$ | H_b, $\epsilon = 0.5$

Figure: Non-private counts and differentially private 2D counts.

Validation on stability of various $\mathcal{L}$

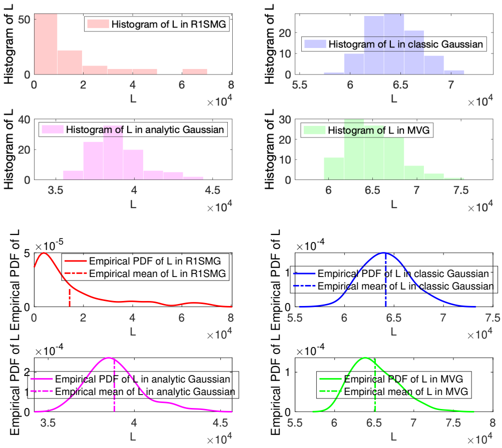

Figure: Accuracy loss introduced by different mechanisms when $\delta = 10^{-7}$, $\epsilon = 10^{-5}$ for the R1SMG mechanism and $\epsilon = 0.5$ for the other mechanisms.

Outline

- Background
- The curse of full-rank covariance matrices
- A new mechanism (R1SMG)
- Open problems

Open problems

- An exact privacy regime
- The impact of the degree of freedom in the noise magnitude on utility, privacy, and stability
- $\frac{1}{sin^2(\theta)} \sim \mathrm{BetaPrime}$. Measure concentration on $\mathrm{BetaPrime}$ r.v. can be leveraged to analyze cumulative privacy loss

Conclusions

- Identify the curse (bottleneck) of utility improvement in existing Gaussian mechanisms
- Propose a new DP mechanism that lifts the curse of full-rank noise covariance matrix
- Leverage measure concentration of random geometric object to bound privacy loss, achieve high utility and stability

Contact : *Tianxi Ji* *tiji@ttu.edu*

Pan Li *lipan@case.edu*