Less is More: Revisiting the Gaussian Mechanism for Differential Privacy

Tianxi Ji, *Texas Tech University;* Pan Li, *Case Western Reserve University*

https://www.usenix.org/conference/usenixsecurity24/presentation/ji

This paper is included in the Proceedings of the 33rd USENIX Security Symposium.

August 14–16, 2024 • Philadelphia, PA, USA

978-1-939133-44-1

Open access to the Proceedings of the 33rd USENIX Security Symposium is sponsored by USENIX.

Less is More: Revisiting the Gaussian Mechanism for Differential Privacy

Tianxi Ji
Texas Tech University
tiji@ttu.edu

Pan Li
Case Western Reserve University
lipan@case.edu

Abstract

Differential privacy (DP) via output perturbation has been a *de facto* standard for releasing query or computation results on sensitive data. Different variants of the classic Gaussian mechanism have been developed to reduce the magnitude of the noise and improve the utility of sanitized query results. However, we identify that all existing Gaussian mechanisms suffer from **the curse of full-rank covariance matrices**, and hence the expected accuracy losses of these mechanisms equal the trace of the covariance matrix of the noise. Particularly, for query results in $\mathbb{R}^M$ (or $\mathbb{R}^{M\times N}$ in a matrix form), in order to achieve (ε,δ)-DP, the expected accuracy loss of the classic Gaussian mechanism, that of the analytic Gaussian mechanism, and that of the Matrix-Variate Gaussian (MVG) mechanism are lower bounded by $C_C(\Delta_2 f)^2$, $C_A(\Delta_2 f)^2$, and $C_M(\Delta_2 f)^2$, respectively, where $C_C = \frac{2\ln\left(\frac{1.25}{\delta}\right)}{\varepsilon^2}M$, $C_A = \frac{\left(\Phi^{-1}(\delta)\right)^2+\varepsilon}{\varepsilon^2}M$, $C_M = \frac{\frac{5}{4}H_r+\frac{1}{4}H_{r,\frac{1}{2}}}{2\varepsilon}MN$, $\Delta_2 f$ is the l_2 sensitivity of the query function f, $\Phi(\cdot)^{-1}$ is the quantile function of the standard normal distribution, H_r is the rth harmonic number, and $H_{r,\frac{1}{2}}$ is the rth harmonic number of order $\frac{1}{2}$.

To lift this curse, we design a **R**ank-1 **S**ingular **M**ultivariate **G**aussian (R1SMG) mechanism. It achieves (ε,δ)-DP on query results in $\mathbb{R}^M$ by perturbing the results with noise following a singular multivariate Gaussian distribution, whose covariance matrix is a **randomly** generated rank-1 positive semi-definite matrix. In contrast, the classic Gaussian mechanism and its variants all consider **deterministic** full-rank covariance matrices. Our idea is motivated by a clue from Dwork et al.'s seminal work on the classic Gaussian mechanism that has been ignored in the literature: when projecting multivariate Gaussian noise with a full-rank covariance matrix onto a set of orthonormal basis in $\mathbb{R}^M$, only the coefficient of a single basis can contribute to the privacy guarantee.

This paper makes the following technical contributions.

(i) The R1SMG mechanisms achieves (ε,δ)-DP guarantee on query results in $\mathbb{R}^M$, while its expected accuracy loss is lower bounded by $C_R(\Delta_2 f)^2$, where $C_R = \frac{2}{\varepsilon\psi}$ and $\psi = \left(\frac{\delta\Gamma(\frac{M-1}{2})}{\sqrt{\pi}\Gamma(\frac{M}{2})}\right)^{\frac{2}{M-2}}$. We show that C_R is on a lower order of magnitude by at least M or MN compared with C_C, C_A, and C_M. Therefore, the expected accuracy loss of the R1SMG mechanism is on a much lower order compared with that of the classic Gaussian mechanism, of the analytic Gaussian mechanism, and of the MVG mechanism.

(ii) Compared with other mechanisms, the R1SMG mechanism is more stable and less likely to generate noise with large magnitude that overwhelms the query results, because the kurtosis and skewness of the nondeterministic accuracy loss introduced by this mechanism is larger than that introduced by other mechanisms.

1 Introduction

Differential privacy (DP) [15, 17] has been increasingly recognized as the fundamental building block for privacy-preserving database query, sharing, and analysis. In this area, one important privacy guarantee is (ε,δ)-DP. It means that, except with a small probability of δ, the presence or absence of any data record in a dataset cannot change the probability of observing a specific query result of the dataset by more than a multiplicative factor of e^{ε}.

The classic Gaussian mechanism, proposed by Dwork et al. [15], is an essential tool to achieve (ε,δ)-DP. Particularly, when the output of a query function is a high dimensional vector, **the classic Gaussian mechanism will add independent and identically distributed (i.i.d.) noise to each component of the result** (see Definition 3 and [17, p. 261], [15, p. 3]).

This standard practice has spread widely into many fields. For example, a differentially private mechanism returns noisy answers to a set of queries by adding i.i.d. noise to each query result [34]. Privacy-preserving learning employs differentially private stochastic gradient descent through perturbing each component of the gradient with i.i.d. Gaussian noise [3]. Unfortunately, as we will show later, the adoption of the classic Gaussian mechanism greatly hinders the utility (accuracy) of the sanitized query results because the expected accuracy loss

(Definition 1) of the classic Gaussian mechanism is on the order of $M(\Delta_2 f)^2$, where $\Delta_2 f$ is the l_2 sensitivity of the query function and M is the dimension of the query results. We refer to this as the **curse of full-rank covariance matrices**.

Definition 1. ***Accuracy Loss [22, 34].*** *The accuracy loss of a differentially private output perturbation mechanism ($\mathcal{M}$) is*

$$\mathcal{L} = ||\mathcal{M}(f(\boldsymbol{x})) - f(\boldsymbol{x})||_2^2 = ||\mathbf{n}||_2^2, \quad (1)$$

where $\boldsymbol{x}$ denotes a dataset, $f(\boldsymbol{x})$ is the query result on dataset $\boldsymbol{x}$, and $\mathbf{n}$ is the noise vector added to the query result.

We can see from (1) that the accuracy loss is equal to the magnitude of the additive noise. Due to the randomness involved in the noise, $\mathcal{L}$ is nondeterministic. Thus, we investigate the expected value of accuracy loss, i.e., $\mathbb{E}[\mathcal{L}]$.

1.1 The Curse of Full-Rank Covariance Matrices

Here, we provide a high-level description on the identified curse of full-rank covariance matrices, while the details are deferred to Section 3.

Proposition 1. ***The Identified Curse.*** *Let $\boldsymbol{x}$ be a dataset, $f(\boldsymbol{x}) \in \mathbb{R}^M$ the queried results, and $\mathbf{n} \in \mathbb{R}^M$ the perturbation noises introduced by the classic Gaussian mechanism (or its variants, e.g., the analytic Gaussian mechanism [5], and the Matrix-Variate Gaussian (MVG) mechanism [11]), respectively. Then, the expected accuracy loss is equal to the trace of the covariance matrix of the Gaussian noise $\mathbf{n}$, i.e., $\mathbb{E}[\mathcal{L}] = Tr[Cov(\mathbf{n})]$, where $Cov(\mathbf{n})$ denotes the covariance matrix of the noise $\mathbf{n}$.*

Since existing Gaussian mechanisms always introduce noise that has a full-rank covariance matrix, i.e., $rank(Cov(\mathbf{n})) = M$, where $Cov(\mathbf{n})$ is symmetric and positive definite [10], $\mathbb{E}[\mathcal{L}]$ is equivalent to the summation of M positive eigenvalues of $Cov(\mathbf{n})$ and hence generally high. In Section 3, we will prove Proposition 1 when $f(\boldsymbol{x})$ is in a vector form and in a matrix form under different variants of the classic Gaussian mechanism.

1.2 Lifting the Curse: Main Contributions

To lift the identified curse, we resort to the **singular multivariate Gaussian distribution** [14, 27, 29, 39, 42] (Definition 7) with a rank-1 covariance matrix and develop the **R**ank-1 **S**ingular **M**ultivariate **G**aussian (R1SMG) mechanism. Our idea is motivated by an ignored clue in Dwork and Roth's proof for the DP guarantee of the Gaussian mechanism [17]. Specifically, Dwork and Roth proved that by projecting the additive noise $\mathbf{n} \in \mathbb{R}^M$ onto a fixed set of orthonormal basis vectors (e.g., $\boldsymbol{b}_1, \boldsymbol{b}_2, \cdots, \boldsymbol{b}_M$), only the coefficient of a single basis vector (e.g., $\mathbf{n}^T \boldsymbol{b}_1$) contributes to the privacy guarantee (more details are deferred to Section 4). However, this important clue has been ignored in the literature. To the best of our knowledge, our work is the first to discover and leverage this ignored clue.

The main contributions of this paper are summarized as follows.

(1) A new DP Mechanism and Privacy Guarantee (Definition 8 and Theorem 5). We develop the R1SMG mechanism that can achieve (ε, δ)-DP guarantee on high dimensional query results $f(\boldsymbol{x}) \in \mathbb{R}^M$ ($M > 2$). The covariance matrix, denoted as $\mathbf{\Pi}$, of the noise introduced by the R1SMG mechanism is a random rank-1 symmetric positive semi-definite matrix, whose eigenvector is randomly sampled from the unit sphere embedded in $\mathbb{R}^M$. Let σ_* be the eigenvalue[1] of the rank-1 covariance matrix $\mathbf{\Pi}$. We prove that a sufficient condition for the R1SMG mechanism to achieve (ε, δ)-DP is $\sigma_* \geq \frac{2(\Delta_2 f)^2}{\varepsilon \psi}$, where $\Delta_2 f$ is the l_2 sensitivity of the query function f, $\psi = \left(\frac{\delta \Gamma(\frac{M-1}{2})}{\sqrt{\pi}\Gamma(\frac{M}{2})}\right)^{\frac{2}{M-2}}$, and $\Gamma(\cdot)$ is Gamma function. Note that similar to that in the classic Gaussian mechanism where $0 < \varepsilon < 1$, we have $0 < \varepsilon < 1/M$ in the R1SMG mechanism. In other words, The R1SMG mechanism can work well under very strict privacy budget.

(2) Expected accuracy loss on a much lower order than that of existing Gaussian mechanisms. We prove that the expected accuracy loss of $f(\boldsymbol{x})$ caused by the R1SMG mechanism (i.e., $||R1SMG(f(\boldsymbol{x})) - f(\boldsymbol{x})||_2^2$) is lower bounded by $C_R(\Delta_2 f)^2$, where for any fixed feasible ε, $C_R = \frac{2}{\varepsilon \psi}$ has a decreasing trend as M increases and converges to $\frac{2}{\varepsilon}$ as M goes large (Theorem 7). In contrast, the classic Gaussian, analytic Gaussian, and MVG mechanisms result in expected accuracy loss that is lower bounded by $C_C(\Delta_2 f)^2$, $C_A(\Delta_2 f)^2$, and $C_M(\Delta_2 f)^2$, respectively, where $C_C = \frac{2\ln(\frac{1.25}{\delta})}{\varepsilon^2} M$, $C_A = \frac{(\Phi^{-1}(\delta))^2 + \varepsilon}{\varepsilon^2} M$, and $C_M = \frac{\frac{5}{4}H_r + \frac{1}{4}H_{r,\frac{1}{2}}}{2\varepsilon} MN$ (for query results in $\mathbb{R}^{M \times N}$, i.e., in a matrix form). Therefore, by using the R1SMG mechanism, **less is more** in the sense that noise of a much lower order of magnitude is needed compared with that of existing Gaussian mechanisms.

(3) Higher stability of the perturbed results (Theorem 8 and 9). We theoretically show that the accuracy of the perturbed $f(\boldsymbol{x})$ obtained by the R1SMG mechanism is more stable than the other mechanisms. By "stable", we mean that it is less likely for the R1SMG mechanism to generate noise with large magnitude that overwhelms $f(\boldsymbol{x})$.

(4) Applications. We perform three case studies, including differentially private 2D histogram query, principal component analysis, and deep learning, to validate that the R1SMG mechanism can achieve better data utility in various applications by introducing less noise, i.e., less is more.

[1] The eigenvalues and singular values are identical for symmetric positive (semi)-definite real-valued matrices. Thus, we use "eigenvalues" and "singular values" interchangeably for such matrices like $\mathbf{\Pi}$ in this paper.

Mechanisms ($\mathcal{M}$)	Expected accuracy loss $\mathbb{E}_{\mathcal{M}}[\mathcal{L}]$	Stability (Section 6)
The classic Gaussian mechanism [15] (noise decided by variance σ^2 and dimension $f(\boldsymbol{x}) \in \mathbb{R}^M$)	$\mathbb{E}_{classic}[\mathcal{L}] = Tr[\sigma^2 \mathbf{I}_{M\times M}] \geq C_C(\Delta_2 f)^2$, where $C_C = \frac{2\ln(\frac{1.25}{\delta})}{\varepsilon^2} M$. (Section 3.1)	**unstable**, i.e., likely to generate noise with large magnitude
The analytic Gaussian mechanism [5] (noise decided by variance σ_A^2 and dimension $f(\boldsymbol{x}) \in \mathbb{R}^M$)	$\mathbb{E}_{analytic}[\mathcal{L}] = Tr[\sigma_A^2 \mathbf{I}_{M\times M}] \geq C_A(\Delta_2 f)^2$, where $C_A = \frac{\left(\Phi^{-1}(\delta)\right)^2 + \varepsilon}{\varepsilon^2} M$. (Section 3.1)	**unstable**
The MVG mechanism [11] (noise decided by covariance matrices $\boldsymbol{\Sigma}$, $\boldsymbol{\Psi}$, and dimension $f(\boldsymbol{x}) \in \mathbb{R}^{M\times N}$)	$\mathbb{E}_{MVG}[\mathcal{L}] = Tr[\boldsymbol{\Sigma} \otimes \boldsymbol{\Psi}] \geq C_M(\Delta_2 f)^2$, where $C_M = \frac{\frac{5}{4} H_r + \frac{1}{4} H_{r,\frac{1}{2}}}{2\varepsilon} MN$. (Section 3.2)	**unstable**
The R1SMG mechanism (noise decided by a random rank-1 singular covariance matrix with only one nonzero eigenvalue σ_*)	$\mathbb{E}_{R1SMG}[\mathcal{L}] = Tr[\boldsymbol{\Pi}] = \sigma_* \geq C_R(\Delta_2 f)^2$, where $C_R = \frac{2}{\varepsilon\psi}$, $\psi = \left(\frac{\delta\Gamma(\frac{M-1}{2})}{\sqrt{\pi}\Gamma(\frac{M}{2})}\right)^{\frac{2}{M-2}}$, and $C_R = \Theta(\frac{\varepsilon C_C}{M}) = \Theta(\frac{C_A}{M}) = \Theta(\frac{C_M}{MN})$. C_R is on a lower order of magnitude by at least M or MN compared with C_C, C_A, and C_M, and converges to $\frac{2}{\varepsilon}$ as M goes large for any fixed feasible ε. (Section 5.3)	**stable**, i.e., unlikely to generate noise with large magnitude

Table 1: Comparisons between the R1SMG mechanism and the classic Gaussian mechanism and its variants.

Table 1 summarizes and compares our proposed R1SMG mechanism with the classic Gaussian mechanism and its variants from the perspective of expected accuracy loss parameterized by the noise parameters and utility stability.

Roadmap. We review the preliminaries for this study in Section 2. Next, we identify the curse of full-rank covariance matrices in various types of Gaussian mechanisms in Section 3. In Section 4, we revisit Dwork and Roth's proof, and identify a previously ignored clue that motivates our work. After that, we formally develop the R1SMG mechanism in Section 5. In Section 6, we analyze the utility stability achieved by our mechanism and theoretically compare it with other mechanisms. Section 7 presents the case studies. Section 8 discusses the related works. Finally, Section 9 concludes the paper.

2 Preliminaries

In this section, we review (ε, δ)-DP, and revisit the definitions and privacy guarantees of the classic Gaussian mechanism and its variants. Table 2 lists the frequently used notations.

ε and δ	privacy budget of the mechanisms
$\Delta_2 f$	l_2 sensitivity of a query function f
σ^2	variance of noise in the classic Gaussian mechanism
σ_A^2	variance of noise in the analytic Gaussian mechanism
$\boldsymbol{\Sigma}$ and $\boldsymbol{\Psi}$	row-, column-wise covariance matrices in MVG
$\boldsymbol{\Pi}$	singular covariance matrix in the R1SMG mechanism
σ_*	the only nonzero eigenvalue of $\boldsymbol{\Pi}$
$\mathbb{PSD}$	set of positive semi-definitive matrices
$\mathbb{PD}$	set of positive definitive matrices
$\mathcal{U}(\mathbb{V}_{1,M})$	uniform distribution on Stiefel manifold $\mathbb{V}_{1,M}$
$\chi^2(M)$	Chi-squared distribution with M degrees of freedom
$\mathcal{L}$	non-deterministic accuracy loss

Table 2: Frequently used notations in the paper.

Following the convention of [15, 17], a database $\boldsymbol{x}$ is represented by its histogram in a universe $\mathcal{X}$: $\boldsymbol{x} \in \mathbb{N}^{|\mathcal{X}|}$, where each entry $\boldsymbol{x}_i$ is the number of elements in the database $\boldsymbol{x}$ of *type* $i \in \mathcal{X}$ [17, p. 17]. $\boldsymbol{x} \sim \boldsymbol{x}'$ denotes a pair of neighboring databases that differ by one data record, and $\Delta_2 f$ is the l_2 sensitivity of a query function f, i.e., $\Delta_2 f = \sup_{\boldsymbol{x}\sim\boldsymbol{x}'} ||f(\boldsymbol{x}) - f(\boldsymbol{x}')||_2$.

Definition 2. *(ε, δ)-DP [17]. A randomized mechanism $\mathcal{M}$ satisfies (ε, δ)-DP if for any two neighboring datasets, $\boldsymbol{x}, \boldsymbol{x}' \in \mathbb{N}^{|\mathcal{X}|}$, $\varepsilon > 0$ and $0 < \delta < 1$, the following holds*

$$\Pr[\mathcal{M}(D_1) \in S] \leq e^{\varepsilon} \Pr[\mathcal{M}(D_2) \in S] + \delta.$$

Dwork et al. [18] also give an alternative definition of (ε, δ)-DP, i.e., for an outcome s, $\ln\left(\frac{\Pr[\mathcal{M}(\boldsymbol{x})=s]}{\Pr[\mathcal{M}(\boldsymbol{x}')=s]}\right) \leq \varepsilon$ holds with all but δ probability. The log-ratio of the probability densities is called the privacy loss random variable [5, 18] (Definition 6).

Definition 3. ***The classic Gaussian Mechanism.*** *Let $f : \mathbb{N}^{|\mathcal{X}|} \to \mathbb{R}^M$ be an arbitrary M-dimensional function. The Gaussian mechanism with parameter σ adds random Gaussian noise following $\mathcal{N}(0, \sigma^2)$ to each of the M components of the output.*

Theorem 1. ***DP Guarantee of the classic Gaussian mechanism [17, p. 261]*** *Let $\varepsilon \in (0, 1)$ be arbitrary. If $\sigma \geq c\Delta_2 f/\varepsilon$, where $c^2 > 2\ln(\frac{1.25}{\delta})$, then the classic Gaussian Mechanism achieves (ε, δ)-DP.*

The analytic Gaussian mechanism [5] follows the same procedure of the classic Gaussian mechanism (discussed in Definition 3). It improves the classic Gaussian mechanism by directly calibrating the variance (σ_A^2) of the Gaussian noise via solving (2) using binary search. The privacy guarantee of this mechanism is stated as follows.

Theorem 2. ***DP Guarantee of the analytic Gaussian mechanism [5].*** *The analytic Gaussian mechanism achieves (ε, δ)-DP guarantee on the query result $f(\boldsymbol{x})$ if*

$$\Phi\left(\frac{\Delta_2 f}{2\sigma_A} - \frac{\varepsilon\sigma_A}{\Delta_2 f}\right) - e^{\varepsilon}\Phi\left(-\frac{\Delta_2 f}{2\sigma_A} - \frac{\varepsilon\sigma_A}{\Delta_2 f}\right) \leq \delta, \qquad (2)$$

where $\Phi(t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{t} e^{-y^2/2} dy$ is the cumulative distribution function of the standard univariate Gaussian distribution.

The MVG mechanism [11] is an extension of the Gaussian mechanism and focuses on matrix-valued queries.

Definition 4. ***The MVG mechanism [11].*** *Given a matrix-valued query function $f(\boldsymbol{x}) \in \mathbb{R}^{M\times N}$, the MVG mechanism is defined as $\mathcal{M}_{MVG}(f(\boldsymbol{x})) = f(\boldsymbol{x}) + \mathbf{N}$, where $\mathbf{N} \sim \mathcal{N}_{M,N}(\mathbf{0}, \boldsymbol{\Sigma}, \boldsymbol{\Psi})$ represents the matrix-variate Gaussian distribution with PDF*

$$f(\mathbf{X}) = \frac{\exp\Big(-\frac{1}{2}Tr[\boldsymbol{\Psi}^{-1}(\mathbf{X}-\mathbf{M})^T\boldsymbol{\Sigma}^{-1}(\mathbf{X}-\mathbf{M})]\Big)}{(2\pi)^{MN/2}\det(\boldsymbol{\Psi})^{M/2}\det(\boldsymbol{\Sigma})^{N/2}}. \quad (3)$$

$\mathbf{M} \in \mathcal{R}^{M\times N}$ is the mean. $\boldsymbol{\Sigma} \in \mathbb{PD}^{M\times M}$ and $\boldsymbol{\Psi} \in \mathbb{PD}^{N\times N}$ are the row-wise and column-wise full-rank covariance matrix, which are also symmetric.

Theorem 3. ***DP Guarantee of MVG [11].*** *$\sigma(\boldsymbol{\Sigma}^{-1})$ and $\sigma(\boldsymbol{\Psi}^{-1})$ are the vectors consisting of singular values of $\boldsymbol{\Sigma}^{-1}$ and $\boldsymbol{\Psi}^{-1}$. The MVG mechanism achieves (ε,δ)-DP if $||\sigma(\boldsymbol{\Sigma}^{-1})||_2||\sigma(\boldsymbol{\Psi}^{-1})||_2 \le \big(-\beta+\sqrt{\beta^2+8\alpha\varepsilon}\big)^2/4\alpha^2$, where $\alpha = 2H_r\gamma\Delta_2 f + (H_r + H_{r,\frac{1}{2}})\gamma^2$, $\beta = 2(MN)^{1/4}H_r\Delta_2 f\tau(\delta)$, H_r is the rth harmonic number, $H_{r,\frac{1}{2}}$ is the rth harmonic number of order $\frac{1}{2}$, $\gamma = \sup_{\boldsymbol{x}}||f(\boldsymbol{x})||_F$, and $\tau(\delta) = 2\sqrt{-MN\ln\delta} - 2\ln\delta + MN$.*

In our study, we also identify an approach to improve the MVG mechanism. Please see Appendix C for details.

3 Identifying The Curse

We identify the curse of full-rank covariance matrices by showing that the expected value of the accuracy loss (Definition 1) introduced by the classic Gaussian mechanism and its variants all are on the order of $M(\Delta_2 f)^2$ for query results in $\mathbb{R}^M$ and $MN(\Delta_2 f)^2$ for query results in $\mathbb{R}^{M\times N}$.

3.1 Query Result being a Vector

When $f(\boldsymbol{x}) \in \mathbb{R}^M$, the classic Gaussian mechanism perturbs each element of $f(\boldsymbol{x})$ using i.i.d. Gaussian noise drawn from $\mathcal{N}(0,\sigma^2)$ (cf. Definition 3). Thus, we derive $\mathbb{E}[\mathcal{L}]$ under the framework of the classic Gaussian mechanism as

$$\mathbb{E}_{\text{classic}}[\mathcal{L}] = \mathbb{E}_{\mathbf{n}_i\sim\mathcal{N}(0,\sigma^2)}\Big[\sum_{i=1}^M \mathbf{n}_i^2\Big] = \mathbb{E}_{z_i\sim\mathcal{N}(0,1)}\Big[\sigma^2\sum_{i=1}^M z_i^2\Big]$$
$$\overset{(a)}{=} \sigma^2\mathbb{E}_{U\sim\chi^2(M)}[U] \overset{(b)}{=} \sigma^2 M = Tr[\sigma^2\mathbf{I}_{M\times M}] = Tr[Cov(\mathbf{n})],$$

where (a) and (b) hold because the squared sum of M independent standard Gaussian variables follows the chi-squared distribution with M degrees of freedom, i.e., $\chi^2(M)$, whose expected value is M.

Based on Theorem 1, since the classic Gaussian mechanism has $\sigma \ge c\Delta_2 f/\varepsilon$, we get $\mathbb{E}_{classic}[\mathcal{L}] \ge C_C(\Delta_2 f)^2$, where $C_C = \frac{2\ln(1.25/\delta)}{\varepsilon^2}M$.

The same analysis also applies to the analytic Gaussian mechanism (discussed in Section 2), because it still perturbs each element of $f(\boldsymbol{x})$ using i.i.d. Gaussian noise with variance σ_A^2. Thus, we can have $\mathbb{E}_{Analytic}[\mathcal{L}] = Tr[\sigma_A^2\mathbf{I}_{M\times M}] = \sigma_A^2 M$.

Since there is no closed-form solution of σ_A ([5] solves for σ_A using a binary search scheme), we derive the lower bound of $\mathbb{E}_{analytic}[\mathcal{L}]$ by investigating the sufficient condition for (2). In what follows $A \Leftarrow B$ means B is the sufficient condition for A, and $A \Leftrightarrow B$ means A and B are equivalent. Then, we have

$$(2) \overset{\text{suff. cond.}}{\Leftarrow} \Phi\Big(\frac{\Delta_2 f}{2\sigma_A} - \frac{\varepsilon\sigma_A}{\Delta_2 f}\Big) \le \delta \Leftrightarrow \frac{\Delta_2 f}{2\sigma_A} - \frac{\varepsilon\sigma_A}{\Delta_2 f} \le \Phi^{-1}(\delta)$$
$$\Leftrightarrow 2\varepsilon\sigma_A^2 + 2\Delta_2 f\Phi^{-1}(\delta)\sigma_A - (\Delta_2 f)^2 \ge 0$$
$$\Leftrightarrow \sigma_A \ge \frac{-2\Delta_2 f\Phi^{-1}(\delta) + \sqrt{4(\Delta_2 f)^2(\Phi^{-1}(\delta))^2 + 8\varepsilon(\Delta_2 f)^2}}{4\varepsilon}$$
$$\overset{\text{suff. cond.}}{\Leftarrow} \sigma_A^2 \ge \frac{\Big(-2\Delta_2 f\Phi^{-1}(\delta) + \sqrt{4(\Delta_2 f)^2(\Phi^{-1}(\delta))^2 + 8\varepsilon(\Delta_2 f)^2}\Big)^2}{16\varepsilon^2}$$
$$\overset{\text{suff. cond.}}{\Leftarrow} \sigma_A^2 \ge \frac{2\big(4(\Delta_2 f)^2(\Phi^{-1}(\delta))^2 + 4(\Delta_2 f)^2(\Phi^{-1}(\delta))^2 + 8\varepsilon(\Delta_2 f)^2\big)}{16\varepsilon^2}$$
$$\Leftrightarrow \sigma_A^2 \ge \frac{\big(\Phi^{-1}(\delta)\big)^2 + \varepsilon}{\varepsilon^2}(\Delta_2 f)^2.$$

Therefore, we can get $\mathbb{E}_{analytic}[\mathcal{L}] \ge C_A(\Delta_2 f)^2$, where $C_A = \frac{\big(\Phi^{-1}(\delta)\big)^2+\varepsilon}{\varepsilon^2}M$.

3.2 Query Result being a Matrix

When $f(\boldsymbol{x}) \in \mathbb{R}^{M\times N}$, instead of adding i.i.d. Gaussian noise, the MVG mechanism by Chanyaswad et al. [11] perturbs the matrix-valued result by adding a matrix noise attributed to the matrix-variate Gaussian distribution shown in (3).

It is well-known that the vectorization of the matrix-variate Gaussian noise can equivalently be sampled from a multivariate Gaussian distribution with a new full-rank covariance matrix, i.e., $\boldsymbol{\Xi} = \boldsymbol{\Sigma}\otimes\boldsymbol{\Psi} \in \mathbb{PD}^{MN\times MN}$, which is also symmetric [21] ($\otimes$ is the Kronecker product). In general, $\boldsymbol{\Xi}$ has nonzero off-diagonal entries, which means the elements in the additive noise are not mutually independent. Thus, we need another tool to analyze its expected accuracy loss.

First, we observe that $\mathcal{L} = ||\mathbf{n}||_2^2 = \mathbf{n}^T\mathbf{n}$ is a quadratic form in random variable defined as follows.

Definition 5. ***Quadratic Form in Random Variable [37].*** *Denote by $\mathbf{x}$ a random vector with mean μ and covariance matrix $\boldsymbol{\Xi}$. The quadratic form in random variable $\mathbf{x}$ associated with a symmetric matrix $\mathbf{A}$ is defined as $Q(\mathbf{x}) = \mathbf{x}^T\mathbf{A}\mathbf{x}$.*

Then, $\mathcal{L}$ introduced by the MVG mechanism is a quadratic form in multivariate Gaussian random variable attributed to $\mathcal{N}(\mathbf{0},\boldsymbol{\Xi})$ with $\mathbf{A}$ being the identity matrix. In the following lemma, we recall the t-th moment of quadratic form in Gaussian random variable (which will also be used in our future analyses on the kurtosis and skewness of $\mathcal{L}$ in Section 6).

Lemma 1. *[37, p. 53] The t-th moment of a quadratic form in multivariate Gaussian random variable, i.e, $\mathbf{x} \sim \mathcal{N}(\mu,\boldsymbol{\Xi})$ with $\boldsymbol{\Xi}$ being full-rank, is given by $\mathbb{E}[Q(\mathbf{x})^t] = \{\sum_{t_1=0}^{t-1}\binom{t-1}{t_1}g^{(t-1-t_1)} \times \sum_{t_2=0}^{t_1-1}\binom{t_1-1}{t_2}g^{(t_1-1-t_2)} \times \cdots \times \sum_{t_2=0}^{1}\binom{1-1}{t_2}g^{(0)}\}$, where $g^{(k)} = 2^k k!(Tr(\mathbf{A}\boldsymbol{\Xi})^{k+1} + (k+1)\mu^T(\mathbf{A}\boldsymbol{\Xi})^k\mathbf{A}\mu)$, $k \in [0,t-1]$.*

As a result, the instantiation of Proposition 1 under the MVG mechanism is

$$\mathbb{E}_{MVG}[\mathcal{L}] = \underset{\mathbf{N}\sim\mathcal{N}_{M,N}(\mathbf{0},\mathbf{\Sigma},\mathbf{\Psi})}{\mathbb{E}}[||\mathbf{N}||_F^2] = \underset{\mathbf{n}\sim\mathcal{N}(\mathbf{0},\mathbf{\Sigma}\otimes\mathbf{\Psi})}{\mathbb{E}}[||\mathbf{n}||_2^2] \overset{(a)}{=} Tr[\mathbf{\Xi}] = Tr[\mathbf{\Sigma}\otimes\mathbf{\Psi}], \quad (4)$$

where (a) is obtained by applying Lemma 1 with $t=1$, i.e., $\mathbb{E}[Q(\mathbf{x})] = g^{(0)} = Tr[\mathbf{\Xi}]$. Furthermore, we can obtain that

$$\begin{aligned}
&\mathbb{E}_{\text{MVG}}[\mathcal{L}] \overset{(a)}{=} Tr[\mathbf{\Sigma}\otimes\mathbf{\Psi}] \overset{(b)}{=} Tr[\mathbf{\Sigma}]Tr[\mathbf{\Psi}] \\
&\overset{(c)}{=} ||\sigma(\mathbf{\Sigma})||_1||\sigma(\mathbf{\Psi})||_1 \overset{(d)}{\geq} ||\sigma(\mathbf{\Sigma})||_2||\sigma(\mathbf{\Psi})||_2 \\
&\overset{(e)}{\geq} \frac{MN}{||\sigma(\mathbf{\Sigma}^{-1})||_2||\sigma(\mathbf{\Psi}^{-1})||_2} \overset{(f)}{\geq} MN\frac{4\alpha^2}{(-\beta+\sqrt{\beta^2+8\alpha\varepsilon})^2} \\
&\geq \frac{MN4\alpha^2}{2\beta^2+8\alpha\varepsilon-2\beta\sqrt{\beta^2}} = MN\frac{\alpha}{2\varepsilon} \\
&\overset{(g)}{=} MN\frac{2H_r\gamma\Delta_2 f + (H_r+H_{r,\frac{1}{2}})\gamma^2}{2\varepsilon} \overset{(h)}{\geq} \frac{\frac{5}{4}H_r+\frac{1}{4}H_{r,\frac{1}{2}}}{2\varepsilon}MN(\Delta_2 f)^2,
\end{aligned} \quad (5)$$

where (a) follows (4), (b) is due to the property of the Kronecker product [40, 41], (c) is because $\mathbf{\Sigma}, \mathbf{\Psi} \in \mathbb{PD}$ and they are both symmetric matrices, which means all their eigenvalues are positive, (d) is because $||\boldsymbol{y}||_1 \geq ||\boldsymbol{y}||_2$ for any vector $\boldsymbol{y}$, (e) is obtained by applying the harmonic mean-geometric mean inequality (see Appendix A), (f) is due to the privacy guarantee of the MVG mechanism in Theorem 3, and α and β are defined therein, (g) is obtained by substituting α, and finally (h) is because $\gamma = \sup_{\boldsymbol{x}} ||f(\boldsymbol{x})||$ (Theorem 3), and $\Delta_2 f = \sup_{\boldsymbol{x}\sim\boldsymbol{x}'} ||f(\boldsymbol{x}) - f(\boldsymbol{x}')|| \leq \sup_{\boldsymbol{x}} ||f(\boldsymbol{x})|| + \sup_{\boldsymbol{x}'} ||f(\boldsymbol{x}')|| = 2\gamma$, which suggests $\gamma \geq \frac{1}{2}\Delta_2 f$.

As a consequence, we have $\mathbb{E}_{MVG}[\mathcal{L}] \geq C_M(\Delta_2 f)^2$, where $C_M = \frac{\frac{5}{4}H_r+\frac{1}{4}H_{r,\frac{1}{2}}}{2\varepsilon}MN$.

In Theorem 4, we summarize the expected accuracy loss of the classic Gaussian mechanism and its variants.

Theorem 4. *The expected accuracy loss of the classic Gaussian, of the analytic Gaussian, and of the MVG mechanisms are as follows:*

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

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

$$\mathbb{E}_{MVG}[\mathcal{L}] = Tr[\mathbf{\Sigma}\otimes\mathbf{\Psi}] \geq C_M(\Delta_2 f)^2, \quad C_M = \frac{\frac{5}{4}H_r+\frac{1}{4}H_{r,\frac{1}{2}}}{2\varepsilon}MN.$$

Theorem 4 rigorously shows that all existing suffer from the identified curse of full-rank covariance matrices. Another stealthy perspective to understand the identified curse of full-rank covariance matrices is by observing that the PDFs of (3) and the multivariate Gaussian distribution all involve determinants of covariance matrices in the denominators (which should be nonzero). Then, theses covariance matrices should always be full-rank. As a consequence, the curse cannot be removed unless we use a new form of PDF to generate the perturbation noise.

4 An Ignored Clue from Dwork et al. [15–17]

In this section, we recall Dwork and Roth's proof of Gaussian mechanism achieving (ε,δ)-DP [15–17] when the query function returns a M-dimension vector. At the end of their proof, we identify an ignored clue which corroborates that Gaussian noise with rank-1 covariance matrix is sufficient to achieve (ε,δ)-DP and also motivates our work.

Dwork and Roth essentially investigate the upper bound of the **privacy loss random variable** (PLRV) on a pair of neighboring database $\boldsymbol{x}$ and $\boldsymbol{x}'$ defined as follows.

Definition 6. *[18, p. 6] Consider running a randomized mechanism $\mathcal{M}$ on a pair of neighboring dataset $\boldsymbol{x}$ and $\boldsymbol{x}'$. For an outcome $\mathbf{s}$, PLRV on $\mathbf{s}$ is defined as the log-ratio of the probability density when $\mathcal{M}$ is running on each dataset, i.e.,* $\text{PLRV}^{(\mathbf{s})}_{(\mathcal{M}(\boldsymbol{x})||\mathcal{M}(\boldsymbol{x}'))} = \ln\left(\frac{\Pr[\mathcal{M}(\boldsymbol{x})=\mathbf{s}]}{\Pr[\mathcal{M}(\boldsymbol{x}')=\mathbf{s}]}\right)$.

The following derivations are restatements of content from [17, p. 261-265]. $f(\cdot)$ is a query function, i.e., $f: \boldsymbol{x} \in \mathbb{N}^{|\mathcal{X}|} \to \mathbb{R}^M$. We are interested in multivariate Gaussian noise that can obscure the difference $\boldsymbol{v} \triangleq f(\boldsymbol{x}) - f(\boldsymbol{x}')$. To achieve (ε,δ)-DP, it requires that the PLRV associated with the classic Gaussian mechanism (denoted as $\mathcal{G}$), i.e.,

$$\text{PLRV}^{(\mathbf{s})}_{(\mathcal{G}(\boldsymbol{x})||\mathcal{G}(\boldsymbol{x}'))} = \left|\frac{1}{2\sigma^2}\left(||\boldsymbol{n}||^2 - ||\boldsymbol{n}+\boldsymbol{v}||^2\right)\right|, \quad (\text{ [17, p. 265]})$$

is upper bounded by ε with all but δ probability. In $\text{PLRV}^{(\mathbf{s})}_{(\mathcal{G}(\boldsymbol{x})||\mathcal{G}(\boldsymbol{x}'))}$, $\boldsymbol{n}$ is chosen from $\mathcal{N}(0,\mathbf{\Sigma})$, where $\mathbf{\Sigma}$ is a diagonal matrix with entries σ^2.

Then, Dwork and Roth bound $\text{PLRV}^{(\mathbf{s})}_{(\mathcal{G}(\boldsymbol{x})||\mathcal{G}(\boldsymbol{x}'))}$ by observing that the multivariate Gaussian distribution is a spherically symmetric distribution [4]. Thus, when representing the noise $\mathbf{n}$ using any fixed orthonormal basis $\boldsymbol{b}_1, \boldsymbol{b}_2, \cdots, \boldsymbol{b}_m$, i.e, $\mathbf{n} = \sum_{i=1}^M \lambda_i \boldsymbol{b}_i$, the corresponding coefficients are also attributed to the same Gaussian distribution, i.e., $\lambda_i \sim \mathcal{N}(0,\sigma^2), i \in [1,M]$. Furthermore, without loss of generality, Dwork and Roth assume the first component (base) $\boldsymbol{b}_1$ is parallel to $\boldsymbol{v}$ (the difference between $f(\boldsymbol{x})$ and $f(\boldsymbol{x}')$). Consequently, we have

$$\begin{aligned}
\text{PLRV}^{(\mathbf{s})}_{(\mathcal{G}(\boldsymbol{x})||\mathcal{G}(\boldsymbol{x}'))} &= \left|\frac{1}{2\sigma^2}\left(\left|\left|\sum_{i=1}^M \lambda_i\boldsymbol{b}_i\right|\right|^2 - \left|\left|\sum_{i=1}^M \lambda_i\boldsymbol{b}_i + \boldsymbol{v}\right|\right|^2\right)\right| \\
&= \left|\frac{1}{2\sigma^2}\left(||\boldsymbol{v}||^2 + 2\lambda_1||\boldsymbol{v}||\right)\right| \leq \frac{1}{2\sigma^2}\left((\Delta_2 f)^2 + 2\lambda_1\Delta_2 f\right),
\end{aligned} \quad (6)$$

where the second equality holds because $\boldsymbol{b}_1$ and $\boldsymbol{v}$ are orthogonal to $\boldsymbol{b}_i, i \in [2,M]$, and the last inequality is due to the definition of l_2 sensitivity. Since now $\text{PLRV}^{(\mathbf{s})}_{(\mathcal{G}(\boldsymbol{x})||\mathcal{G}(\boldsymbol{x}'))}$ in (6) only involves a single Gaussian random variable, i.e., λ_1, Theorem 1 can be proved by following the same procedure when the query function returns a scalar value ([17, p. 262-264]).

A hidden Clue comes up to the surface. From (6), it is clear that $\text{PLRV}^{(\mathbf{s})}_{(\mathcal{G}(\boldsymbol{x})||\mathcal{G}(\boldsymbol{x}'))}$ is only related to λ_1 and $\Delta_2 f$. Since λ_1 (a Gaussian random variable) is the coefficient of

the of orthonormal base $\boldsymbol{b}_1$, it clearly suggests that, after decomposing a multivariate Gaussian noise using a set of orthonormal basis vectors, only the coefficient of one vector contributes to the value of $\text{PLRV}^{(\mathbf{s})}_{(\mathcal{G}(\boldsymbol{x})||\mathcal{G}(\boldsymbol{x}'))}$. In other words, the coefficients of other basis vectors have no impact on the privacy loss of the Gaussian mechanism. Consequently, the privacy guarantee achieved by $\mathbf{n} = \sum_{i=1}^{M} \lambda_i \boldsymbol{b}_i$ is identical to that achieved by $\lambda_1 \boldsymbol{b}_1$. This hidden clue motivates our idea of using multivariate Gaussian noise whose covariance matrix has rank-1 (instead of multivariate Gaussian noise with full-rank covariance matrix) to achieve (ε, δ)-DP.

5 Lifting the Curse

In what follows, we present the R1SMG mechanism, which lifts the identified curse of full-rank covariance matrices. In particular, we first intuitively explain the feasibility of the proposed R1SMG mechanism. Next, we introduce its noise generation process, and present a sufficient condition for it to achieve (ε, δ)-DP. After that, we analyze the expected accuracy loss of the R1SMG mechanism, and discuss the privacy leakage of utilizing vectors in the null space of the noise.

5.1 The Intuition behind R1SMG

As introduced in Section 1.2, the expected accuracy loss of R1SMG is lower bounded by $C_R(\Delta_2 f)^2$ where for any fixed feasible ε, $C_R = \frac{2}{\varepsilon\psi}$ has a decreasing trend as M increases, and converges to $\frac{2}{\varepsilon}$ as M goes large (Theorem 7). Thus, we have the following property (see Property 1) that is missing in all existing DP mechanisms. Here, we provide an intuitive explanation of it.

Property 1. *For any fixed feasible ε, δ, and $\Delta_2 f$, the magnitude of the noise required to attain (ε, δ)-DP on $f(\boldsymbol{x}) \in \mathbb{R}^M$ can have a non-increasing trend regarding M.*

Due to the widely accepted common practice of perturbing each component of $f(\boldsymbol{x})$ using i.i.d. Gaussian noise to achieve DP, it makes sense that larger dimensional $f(\boldsymbol{x})$ requires larger magnitude of noises. Thus, Property 1 is counterintuitive and seems to be "magic". Yet, it can be intuitively explained as follows. Given a database $\boldsymbol{x}$ represented as histogram, we consider a normalized counting query function $f(\boldsymbol{x}) = \frac{1}{M} Q\boldsymbol{x} \in \mathbb{R}^M$, where Q is a binary matrix. A higher dimension of $f(\boldsymbol{x})$ means a larger number of queries applied to the same dataset $\boldsymbol{x}$. Subsequently, the chances that the query results (i.e., rows in Q) become dependent with each other increases as the query number increases. For instance, suppose that the ith and jth query, i.e., $f(\boldsymbol{x})_i$ and $f(\boldsymbol{x})_j$ are dependent, particularly, $f(\boldsymbol{x})_j$ can be fully determined by $f(\boldsymbol{x})_i$. Then, the privacy leakage on $\boldsymbol{x}$ by observing $f(\boldsymbol{x})_i$, $f(\boldsymbol{x})_j$, or $[f(\boldsymbol{x})_i\ f(\boldsymbol{x})_j]$ would be identical. Hence, intuitively, no more noise is needed to perturb $[f(\boldsymbol{x})_i\ f(\boldsymbol{x})_j]$ than that for perturbing $f(\boldsymbol{x})_i$. It implies that as M increases and the queried results become dependent, it is possible to achieve DP by perturbing $f(\boldsymbol{x})$ using noise of a non-increasing magnitude. Moreover, when M becomes sufficiently large such that the query vectors (rows in Q) are linearly dependent, by observing $f(\boldsymbol{x})$ or the set of independent query vectors, the privacy leakage of $\boldsymbol{x}$ would be identical. Therefore, the same amount of noise can be used to sanitize both (refer to [43] for an analysis from the perspective of mutual information).

5.2 R1SMG: Multivariate Gaussian Noise with A Random Rank-1 Covariance Matrix

In what follows, we first provide the statistical background on the singular multivariate Gaussian distribution with a given rank-r covariance matrix, e.g., $\boldsymbol{\Pi}$ and $Rank(\boldsymbol{\Pi}) = r$.

Definition 7. ***Singular Multivariate Gaussian Distribution [14, 27, 39, 42].*** *A M-dimensional random variable $\mathbf{x} = [x_1, x_2, \cdots, x_M]^T \in \mathbb{R}^M$ has a singular multivariate Gaussian distribution with mean $\mu \in \mathbb{R}^M$ and a singular covariance matrix $\boldsymbol{\Pi} \in \mathbb{PSD}^{M\times M}$ with rank-r, i.e., $\mathbf{x} \sim \mathcal{N}(\mu, \boldsymbol{\Pi})$ and $Rank(\boldsymbol{\Pi}) = r < M$, if its PDF is*

$$f_{\mathbf{x};Rank(\boldsymbol{\Pi})=r} = \frac{(2\pi)^{-r/2} \exp\left(-\frac{1}{2}(\mathbf{x}-\mu)^T \boldsymbol{\Pi}^{\dagger}(\mathbf{x}-\mu)\right)}{\sqrt{\sigma_1(\boldsymbol{\Pi}) \cdots \sigma_r(\boldsymbol{\Pi})}},$$

where $\sigma_i(\boldsymbol{\Pi})$ is the i-th nonzero eigenvalue of $\boldsymbol{\Pi}$, and $\boldsymbol{\Pi}^{\dagger}$ is the Moore-Penrose generalized inverse of $\boldsymbol{\Pi}$. Particularly, we have $\boldsymbol{\Pi}^{\dagger} = \mathbf{V}_r \boldsymbol{\Lambda}_r^{-1} \mathbf{V}_r^T$, where $\boldsymbol{\Lambda}_r = \text{diag}(\sigma_1(\boldsymbol{\Pi}), \ldots, \sigma_r(\boldsymbol{\Pi})) \in \mathbb{R}^{r\times r}$, and $\mathbf{V}_r \in \mathbb{R}^{M\times r}$ is the matrix of eigenvectors corresponding to the r nonzero eigenvalues.

Samples of $f_{\mathbf{x};Rank(\boldsymbol{\Pi})=r}$ (e.g., $\mathbf{n}$) can be generated by applying the disintegration theorem [20]. In particular, if $\mu = \mathbf{0}$, by defining a linear mapping

$$\mathbf{n} = \mathbf{V}_r \boldsymbol{\Lambda}_r^{1/2} \mathbf{z}, \quad \text{and} \quad \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{r\times r}), \tag{7}$$

where $\boldsymbol{\Lambda}_r^{1/2} = \text{diag}(\sqrt{\sigma_1(\boldsymbol{\Pi})}, \sqrt{\sigma_2(\boldsymbol{\Pi})}, \ldots, \sqrt{\sigma_r(\boldsymbol{\Pi})}) \in \mathbb{R}^{r\times r}$ and $\mathbf{V}_r \in \mathbb{R}^{M\times r}$ is the matrix of eigenvectors corresponding to the r nonzero eigenvalues of $\boldsymbol{\Pi}$, then $\mathbf{n}$ is attributed to a singular multivariate Gaussian distribution with a rank-r covariance matrix. This can easily be verified by checking the covariance matrix of $\mathbf{n}$;

$$\begin{aligned} Cov(\mathbf{n}) =& Cov(\mathbf{V}_r \boldsymbol{\Lambda}_r^{1/2} \mathbf{z}) = \mathbf{V}_r \boldsymbol{\Lambda}_r^{1/2} Cov(\mathbf{z}) (\mathbf{V}_r \boldsymbol{\Lambda}_r^{1/2})^T \\ \overset{(*)}{=}& \mathbf{V}_r \boldsymbol{\Lambda}_r^{1/2} \mathbf{I}_{r\times r} (\mathbf{V}_r \boldsymbol{\Lambda}_r^{1/2})^T = \mathbf{V}_r \boldsymbol{\Lambda}_r \mathbf{V}_r^T = \boldsymbol{\Pi} \end{aligned}, \tag{8}$$

where $(*)$ is due to $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{r\times r})$. Then, clearly, we have $Cov(\mathbf{n}) = r$.

When $r = 1$, the singular covariance matrix $\boldsymbol{\Pi}$ only has 1 nonzero eigenvalue, and we denote it as σ_*. Thus, (7) becomes

$$\mathbf{n} = \mathbf{v}\sqrt{\sigma_*} z, \quad \text{and} \quad z \sim \mathcal{N}(0, 1), \quad \mathbf{v}^T \mathbf{v} = 1. \tag{9}$$

Singular multivariate Gaussian noise with a random rank-1 covariance matrix. It is noteworthy that in (9), $\mathbf{v}$ needs to be generated randomly to thwart attacks which takes advantage of vectors in the null space of $\mathbf{v}$ (or $\mathbf{n}$). In Section 5.4, we will show that by sampling $\mathbf{v}$ uniformly at random, such attacks will succeed with **zero probability**.

Since $\mathbf{v}$ is an orthonormal vector of dimension $M \times 1$, one common approach to generate $\mathbf{v}$ randomly is by uniformly sampling from a specific Stiefel manifold, i.e., $\mathbb{V}_{1,M} = \{\mathbf{v} \in \mathbb{R}^{M\times 1} : \mathbf{v}^T\mathbf{v} = 1\}$, which represents the unite sphere $\mathbb{S}^{M-1}$ embedded in $\mathbb{R}^M$. In Statistics literature, the PDF of the uniform distribution on the Stiefel manifold $\mathbb{V}_{1,M}$ is given by

$$f(\mathbf{v}) = \frac{1}{\mathrm{Vol}(\mathbb{V}_{1,M})}, \quad \forall \mathbf{v} \in \mathbb{V}_{1,M}, \tag{10}$$

(see [21, p. 280], [8, p. 17, equation 8.2.2], as well [13, p. 30] which gives the characteristic function of (10)). In particular, the constant $\mathrm{Vol}(\mathbb{V}_{1,M}) = \frac{2\pi^{M/2}}{\Gamma(M/2)}$ (see [21, p. 19 and 26]) is the total surface area or volume of $\mathbb{V}_{1,M}$ and $\Gamma(\cdot)$ is the ordinary gamma function [21][2].

We use $\mathbf{v} \sim \mathcal{U}(\mathbb{V}_{1,M})$ to represent that $\mathbf{v}$ is a random variable uniformly sampled from $\mathbb{V}_{1,M}$. As a consequence, the linear mapping defined in (9) becomes

$$\mathbf{n} = \mathbf{v}\sqrt{\sigma_*}z, \quad \text{where} \quad z \sim \mathcal{N}(0,1), \quad \mathbf{v} \sim \mathcal{U}(\mathbb{V}_{1,M}). \tag{11}$$

Now, we introduce the R1SMG mechanism and provide a sufficient condition for it to achieve differential privacy.

Definition 8. ***The Rank-1 Singular Multivariate Gaussian (R1SMG) Mechanism.*** *For an arbitrary M-dimensional query function, $f(\boldsymbol{x}) \in \mathbb{R}^M$, the R1SMG mechanism is defined as*

$$\mathcal{M}_{R1SMG}(f(\boldsymbol{x})) = f(\boldsymbol{x}) + \mathbf{n},$$

where $\mathbf{n}$ (generated via (11)) is the noise attributed to a singular multivariate Gaussian distribution with a random rank-1 covariance matrix.

If the query result is a matrix or tensor, we can first generate the noise as a vector and then resize it into the desired format. Please refer to the case studies in Section 7 for details.

We introduce an important lemma below that will be used for proving Theorem 5.

Lemma 2. ***Distribution of Angle between Random Points on Unit Sphere ([7, p. 1845 and 1860]).*** *Let h and g be two randomly selected points on unit sphere $\mathbb{S}^{P-1}$ (embedded in $\mathbb{R}^P$), where $P > 2$. Let $\mathbf{h}$ (resp. $\mathbf{g}$) be the vector connecting the center of the sphere and h (resp. g). θ denotes the angle ($\theta \in [0,\pi]$) between $\mathbf{h}$ and $\mathbf{g}$. Then, we have*

$$\Pr\left[\left|\theta - \frac{\pi}{2}\right| \geq \theta_0\right] \leq \sqrt{\pi}\frac{\Gamma(\frac{P}{2})}{\Gamma(\frac{P-1}{2})} cos(\theta_0)^{P-2}, \tag{12}$$

where θ_0 is a given radian and $0 \leq \theta_0 \leq \frac{\pi}{2}$.

[2]This result can be easily obtained by setting $p = 1$ in Theorem 1.4.9 on page 25 of [21]. In particular, $\mathrm{Vol}(\mathbb{V}_{1,M})$ plays the same role on $\mathbb{V}_{1,M}$ as the Lebesgue measure plays in Euclidean space [8]. For example, when $M = 3$, $\mathrm{Vol}(\mathbb{V}_{1,3}) = 4\pi$ is the surface area of an unit sphere in 3D space, and when $M = 2$, $\mathrm{Vol}(\mathbb{V}_{1,2}) = 2\pi$ is the circumference of an unit circle in 2D space.

Theorem 5. *The R1SMG mechanism achieves (ε,δ)-DP when $M > 2$, if $\sigma_* \geq \frac{2(\Delta_2 f)^2}{\varepsilon\psi}$ 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. Inspired by Dwork and Roth's work [17] (cf. Section 4), we investigate the PLRV associated with the R1SMG mechanism. Assume that $\mathbf{n}$ and $\mathbf{n}'$ are the random noise used to perturb $f(\boldsymbol{x})$ and $f(\boldsymbol{x}')$, respectively. Then, the R1SMG mechanism achieves (ε,δ)-DP if $\mathrm{PLRV}^{(\mathbf{s})}_{(R1SMG(\boldsymbol{x})||R1SMG(\boldsymbol{x}'))} = \ln\left(\frac{\Pr[f(\boldsymbol{x})+\mathbf{n}=\mathbf{s}\in\mathcal{S}]}{\Pr[f(\boldsymbol{x}')+\mathbf{n}'=\mathbf{s}\in\mathcal{S}]}\right) \leq \varepsilon$ with all but δ probability (called the failing probability) [17], where $\mathcal{S}$ denotes all possible outcome of R1SMG.

According to (11), without loss of generality, let $\mathbf{n} = \sqrt{\sigma_*}z_1\mathbf{h}$ and $\mathbf{n}' = \sqrt{\sigma_*}z_2\mathbf{g}$, where $z_1, z_2 \sim \mathcal{N}(0,1)$ and $\mathbf{h},\mathbf{g} \sim \mathcal{U}(\mathbb{V}_{1,M})$. Let $\theta \in [0,\pi]$ be the angle between $\mathbf{h}$ and $\mathbf{g}$. Then, we can establish the failing probability (which should be at most δ) by first constructing an event which is a subspace of the universe Ω, i.e.,

$$\begin{aligned}\mathcal{A} &= \left\{z_1,z_2,\mathbf{h},\mathbf{g} : z_1,z_2 \sim \mathcal{N}(0,1), \mathbf{h},\mathbf{g} \sim \mathcal{U}(\mathbb{V}_{1,M}), \left|\theta - \tfrac{\pi}{2}\right| < \theta_0\right\} \\ &\subset \Omega = \left\{z_1,z_2,\mathbf{h},\mathbf{g} : z_1,z_2 \sim \mathcal{N}(0,1), \mathbf{h},\mathbf{g} \sim \mathcal{U}(\mathbb{V}_{1,M})\right\}\end{aligned},$$

where $\theta_0 = arccos\left(\left(\delta\Gamma(\frac{M-1}{2})/\sqrt{\pi}\Gamma(\frac{M}{2})\right)^{1/(M-2)}\right)$, and $arccos(\cdot)$ is the inverse of the cosine function. We denote the complementary of $\mathcal{A}$ as $\mathcal{A}^c = \Omega \setminus \mathcal{A}$, i.e.,

$$\mathcal{A}^c = \left\{z_1,z_2,\mathbf{h},\mathbf{g} : z_1,z_2 \sim \mathcal{N}(0,1), \mathbf{h},\mathbf{g} \sim \mathcal{U}(\mathbb{V}_{1,M}), \left|\theta - \tfrac{\pi}{2}\right| \geq \theta_0\right\}.$$

Since z_1 and z_2 are independent of $\mathbf{h}$ and $\mathbf{g}$, by applying Lemma 2 and setting $P = M$ and $\sqrt{\pi}\frac{\Gamma(\frac{P}{2})}{\Gamma(\frac{P-1}{2})} cos(\theta_0)^{P-2} = \delta$, we get $\Pr[\mathcal{A}^c] \leq \delta$.

Since $\mathbf{n}$ (resp. $\mathbf{n}'$) is singular multivariate Gaussian with zero mean by design (see (11)), we have that $f(\boldsymbol{x}) + \mathbf{n} = \mathbf{s} \in \mathcal{S}$ (resp. $f(\boldsymbol{x}') + \mathbf{n}' = \mathbf{s} \in \mathcal{S}$) is attributed to singular multivariate Gaussian distribution, with mean $f(\boldsymbol{x})$ (resp. $f(\boldsymbol{x}')$) and covariance matrix $\mathbf{h}\sigma_*\mathbf{h}^T$ (resp. $\mathbf{g}\sigma_*\mathbf{g}^T$) (one can verify this by setting $r = 1$ in (8)). Substituting the mean and generalized inverse covariance matrix into the PDF provided in Definition 7 (with $r = 1$), we have a counterpart of (6) as

$$\begin{aligned}&\mathrm{PLRV}^{(\mathbf{s})}_{(R1SMG(\boldsymbol{x})||R1SMG(\boldsymbol{x}'))} = \ln\left(\frac{\Pr[f(\boldsymbol{x})+\mathbf{n}=\mathbf{s}\in\mathcal{S}]}{\Pr[f(\boldsymbol{x}')+\mathbf{n}'=\mathbf{s}\in\mathcal{S}]}\right) \\ =&\ln\left(\frac{(2\pi)^{-\frac{1}{2}}\exp\left(-\frac{1}{2}\left(\mathbf{s}-f(\boldsymbol{x})\right)^T\left(\mathbf{h}\sigma_*^{-1}\mathbf{h}^T\right)\left(\mathbf{s}-f(\boldsymbol{x})\right)\right)\Big/\sqrt{\sigma_*}}{(2\pi)^{-\frac{1}{2}}\exp\left(-\frac{1}{2}\left(\mathbf{s}-f(\boldsymbol{x}')\right)^T\left(\mathbf{g}\sigma_*^{-1}\mathbf{g}^T\right)\left(\mathbf{s}-f(\boldsymbol{x}')\right)\right)\Big/\sqrt{\sigma_*}}\right) \\ =&\ln\left(\frac{\exp\left(-\frac{1}{2\sigma_*}\left(\mathbf{h}^T(\mathbf{s}-f(\boldsymbol{x}))\right)^2\right)}{\exp\left(-\frac{1}{2\sigma_*}\left(\mathbf{g}^T(\mathbf{s}-f(\boldsymbol{x}'))\right)^2\right)}\right) \\ =&\frac{1}{2\sigma_*}\left(\left(\underbrace{\mathbf{g}^T(\mathbf{s}-f(\boldsymbol{x}'))}_{\rho_1\in\mathbb{R}}\right)^2 - \left(\underbrace{\mathbf{h}^T(\mathbf{s}-f(\boldsymbol{x}))}_{\rho_2\in\mathbb{R}}\right)^2\right) \\ \leq&\frac{1}{2\sigma_*}\left(|\rho_1|+|\rho_2|\right)^2, \qquad \forall(z_1,z_2,\mathbf{h},\mathbf{g}) \in \mathcal{A}.\end{aligned} \tag{13}$$

In the following, we derive a tight upper bound on $|\rho_1|+|\rho_2|$. In particular, with $||\mathbf{g}||_2 = ||\mathbf{h}||_2 = 1$, we notice that

$$|\rho_1|+|\rho_2| \leq ||\mathbf{g}||_2||\mathbf{s}-f(\boldsymbol{x}')||_2 + ||\mathbf{h}||_2||\mathbf{s}-f(\boldsymbol{x})||_2$$
$$= ||\mathbf{s}-f(\boldsymbol{x}')||_2 + ||\mathbf{s}-f(\boldsymbol{x})||_2.$$

Thus, we aim to bound $||\mathbf{s}-f(\boldsymbol{x}')||_2 + ||\mathbf{s}-f(\boldsymbol{x})||_2$.

First, we observe that the angle between $(\mathbf{s}-f(\boldsymbol{x}))$ and $(\mathbf{s}-f(\boldsymbol{x}'))$ is identical to the angle between $\mathbf{h}$ and $\mathbf{g}$, because $(\mathbf{s}-f(\boldsymbol{x})) = \mathbf{n}$ and $(\mathbf{s}-f(\boldsymbol{x}')) = \mathbf{n}'$, and $\mathbf{n}$ and $\mathbf{n}'$ are obtained by scaling $\mathbf{h}$ and $\mathbf{g}$. We use θ to represent the angle between $\mathbf{h}$ and $\mathbf{g}$ (i.e., identically between $(\mathbf{s}-f(\boldsymbol{x}))$ and $(\mathbf{s}-f(\boldsymbol{x}'))$). Thus, from a geometric perspective, $||\mathbf{s}-f(\boldsymbol{x}')||_2 + ||\mathbf{s}-f(\boldsymbol{x})||_2$ is the sum of the length of two edges of a triangle in $\mathbb{R}^M$, and the angle between these two edges is θ. Moreover, the length of the third edge (opposite to this specific angle) is $||(\mathbf{s}-f(\boldsymbol{x}'))-(\mathbf{s}-f(\boldsymbol{x}))||_2 = ||f(\boldsymbol{x})-f(\boldsymbol{x}')||_2$. For better understanding, we visualize the above description in Figure 1(a), where the red point (resp. blue point) indicates $\mathbf{n}$ (resp. $\mathbf{n}'$) in $\mathbb{R}^M$, and the red dashed line (resp. blue dashed line) is the randomly sampled $\mathbf{h}$ (resp. $\mathbf{g}$) (note that we show $\mathbf{h}$ and $\mathbf{g}$ as bidirectional, because $\mathbf{n}$ and $\mathbf{n}'$ can point to the opposite direction of $\mathbf{h}$ and $\mathbf{g}$, respectively).

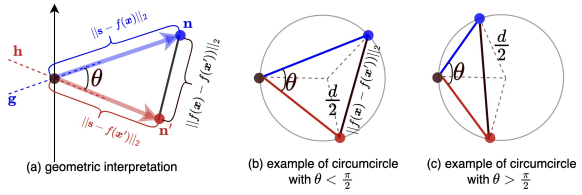

Figure 1: (a) Geometric interpretation of $|\rho_1|+|\rho_2|$. Circumcircles of the described triangle with (b) $\theta < \frac{\pi}{2}$ and (c) $\theta > \frac{\pi}{2}$.

It is well-known that the length of each edge of a triangle is upper bounded by the diameter of the circumscribed circle of the triangle. We show this fact in Figure 1(b) and Figure (c), where $\frac{d}{2}$ represents the radius of the circumcircle. According to the law of sine, we have $d = ||f(\boldsymbol{x})-f(\boldsymbol{x}')||_2 / sin(\theta)$. Thus,

$$||\mathbf{s}-f(\boldsymbol{x}')||_2 + ||\mathbf{s}-f(\boldsymbol{x})||_2 \leq 2d = \frac{2||f(\boldsymbol{x})-f(\boldsymbol{x}')||_2}{sin(\theta)} \leq \frac{2\Delta_2 f}{sin(\theta)}, \quad (14)$$

where the last inequality is because the query sensitivity on neighboring datasets is $\Delta_2 f = \sup_{\boldsymbol{x}\sim\boldsymbol{x}'} ||f(\boldsymbol{x})-f(\boldsymbol{x}')||_2$.

In (13), we have $(z_1, z_2, \mathbf{h}, \mathbf{g}) \in \mathcal{A}$, i.e., $\theta \in (\frac{\pi}{2}-\theta_0, \frac{\pi}{2}+\theta_0)$, $\theta_0 \in [0, \frac{\pi}{2}]$. Thus, due to the symmetry of the sine function around $\frac{\pi}{2}$, we can obtain $\frac{2\Delta_2 f}{sin(\theta)} \leq \frac{2\Delta_2 f}{sin(\frac{\pi}{2}-\theta_0)} = \frac{2\Delta_2 f}{cos(\theta_0)}$. As a result, we get $|\rho_1|+|\rho_2| \leq ||\mathbf{s}-f(\boldsymbol{x}')||_2 + ||\mathbf{s}-f(\boldsymbol{x})||_2 < \frac{2\Delta_2 f}{cos(\theta_0)} = 2\Delta_2 f \Big/ \left(\frac{\delta\Gamma(\frac{M-1}{2})}{\sqrt{\pi}\Gamma(\frac{M}{2})}\right)^{\frac{1}{M-2}}$.

Finally, from (13), we can obtain the following sufficient condition for the R1SMG mechanism to achieve (ε, δ)-DP:

$$\mathrm{PLRV}^{(\mathbf{s})}_{(R1SMG(\boldsymbol{x})||R1SMG(\boldsymbol{x}'))}$$
$$\leq \frac{1}{2\sigma_*}(|\rho_1|+|\rho_2|)^2 \leq \frac{1}{2\sigma_*}\left(2\Delta_2 f \Big/ \left(\frac{\delta\Gamma(\frac{M-1}{2})}{\sqrt{\pi}\Gamma(\frac{M}{2})}\right)^{\frac{1}{M-2}}\right)^2 \leq \varepsilon,$$

which leads to $\sigma_* \geq \frac{2(\Delta_2 f)^2}{\varepsilon\psi}$. Moreover, due to the constraint of $P > 2$ in Lemma 2, we require $M > 2$ for the R1SMG mechanism. Thus, Theorem 5 follows. □

Remark 1. *In the classic Gaussian mechanism, the privacy budget is upper bounded by 1 (Theorem 1). Similarly, there is also an upper bound on ε for the R1SMG mechanism. This can be explained geometrically using Figure 2. In particular, Figure 2(a) shows the case where $\varepsilon < 1$ in the classic Gaussian mechanism. Clearly as long as $\boldsymbol{n}$ and $\boldsymbol{n}'$ can obscure the difference between $f(\boldsymbol{x})$ and $f(\boldsymbol{x})'$, i.e., $f(\boldsymbol{x})+\boldsymbol{n} = f(\boldsymbol{x}')+\boldsymbol{n}' = \mathbf{s}$, the noise components along the direction of $\boldsymbol{v} = f(\boldsymbol{x})-f(\boldsymbol{x}')$ are also sufficient to obscure the difference between $f(\boldsymbol{x})$ and $f(\boldsymbol{x})'$, i.e., $\lambda_1\boldsymbol{b}_1 - (\boldsymbol{b}_1^T\boldsymbol{n}')\boldsymbol{b}_1 = \boldsymbol{v}$. Whereas, if ε exceeds the upper bound, the magnitude of $\boldsymbol{n}$ and $\boldsymbol{n}'$ might be too small to obscure $\boldsymbol{v}$, since $||\boldsymbol{n}||$ and $||\boldsymbol{n}'||$ are proportional to $\frac{1}{\varepsilon}$. In other words, when ε is beyond the upper bound, $\boldsymbol{n}$, $\boldsymbol{n}'$, and $\boldsymbol{v}$ cannot form a triangle, and hence we have $\lambda_1\boldsymbol{b}_1 - (\boldsymbol{b}_1^T\boldsymbol{n}')\boldsymbol{b}_1 \neq \boldsymbol{v}$ as shown in Figure 2(b). This is also true for the R1SMG mechanism, i.e., when ε is too large, $\boldsymbol{n} = \sqrt{\sigma_*}z_1\mathbf{h}$, $\boldsymbol{n}' = \sqrt{\sigma_*}z_2\mathbf{g}$, and $\boldsymbol{v}$ cannot form a triangle as shown in Figure 2(c). As a result, in order to form a triangle as shown in Figure 2(d), an upper bounded ε is required for the R1SMG mechanism. In this work, we let the privacy budget of the R1SMG mechanism be upper bounded by $\frac{1}{M}$ of that of the classic Gaussian mechanism, i.e., $\frac{1}{M}$. This is because $\boldsymbol{v}$ is obscured using noise with M degrees of freedom by the classic Gaussian mechanism, whereas it is obscured using noise with 1 degree of freedom by the R1SMG mechanism due to the rank-1 constraint. We will provide a tight upper bound on ε for the R1SMG mechanism in a separate study.*

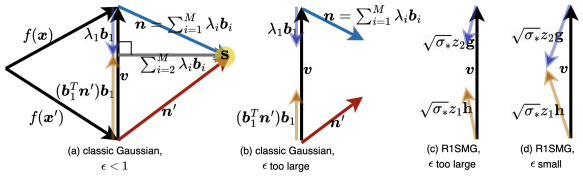

Figure 2: Geometric interpretations on the constraints on ε in the classic Gaussian and R1SMG mechanism.

Implementation of R1SMG. According to the following theorem, we can obtain the desired random variable distributed on the Stiefel manifold by transforming samples drawn i.i.d. from standard Gaussian distribution.

Theorem 6. *[13, p. 29, Theorem 2.2.1] Let $\mathbf{x} \in \mathbb{R}^{M\times r}$, whose elements are i.i.d. Gaussian random variables from $\mathcal{N}(0,1)$. Then, $\mathbf{v} = \mathbf{x}(\mathbf{x}^T\mathbf{x})^{-1/2}$ is uniformly distributed on $\mathbb{V}_{r,M}$.*

Thus, we can apply Theorem 6 with $r = 1$ to draw the desired samples from $\mathbb{V}_{1,M}$. Then, applying (11), we can generate an instance of random noise for the R1SMG mechanism.

5.3 Expected Accuracy Loss Analysis

In this section, we investigate the expected accuracy loss of the R1SMG mechanism, i.e., $\mathbb{E}_{R1SMG}[\mathcal{L}]$. The results are summarized in Theorem 7.

Theorem 7. *For any fixed feasible $\varepsilon > 0, 0 < \delta < 1$, given a dataset $\boldsymbol{x}$ and a query result $f(\boldsymbol{x}) \in \mathbb{R}^M$. We have $\mathbb{E}_{R1SMG}[\mathcal{L}] = ||R1SMG(f(\boldsymbol{x})) - f(\boldsymbol{x})||_2^2 = Tr[\boldsymbol{\Pi}] = \sigma_* \geq C_R(\Delta_2 f)^2$, where σ_* is the only nonzero eigenvalue of $\boldsymbol{\Pi}$, $C_R = \frac{2}{\varepsilon\psi}$, and $\psi = \left(\frac{\delta\Gamma(\frac{M-1}{2})}{\sqrt{\pi}\Gamma(\frac{M}{2})}\right)^{\frac{2}{M-2}}$. C_R has a decreasing trend as M increases. When M goes large, C_R converges to $\frac{2}{\varepsilon}$ and $\mathbb{E}_{R1SMG}[\mathcal{L}]$ can be as low as $\frac{2(\Delta_2 f)^2}{\varepsilon}$.*

Proof. According to the noise generation process in R1SMG, i.e., (11), we have

$$\begin{aligned}&\mathbb{E}_{R1SMG}[\mathcal{L}] = \mathbb{E}[||\mathbf{n}||_2^2]\\ =&\mathbb{E}\left[(\mathbf{v}\sigma_*^{1/2}z)^T\mathbf{v}\sigma_*^{1/2}z\right] = \sigma_*\mathbb{E}[z^2] \overset{(a)}{=} \sigma_* \overset{(b)}{\geq} 2(\Delta_2 f)^2/(\varepsilon\psi),\end{aligned} \tag{15}$$

where (a) holds because $z^2 \sim \chi^2(1)$, and (b) is due to Theorem 5. From (15), we also have $\mathbb{E}_{R1SMG}[\mathcal{L}] = \sigma_* = Tr[\boldsymbol{\Pi}]$.

To show the decreasing trend of $\mathbb{E}_{R1SMG}[\mathcal{L}]$ without dealing with the cumbersome notation of ψ, we use some intermediate results obtained in the proof of Theorem 5. To be more specific, we have $\psi = sin^2(\theta)$, where $\theta \in (\frac{\pi}{2} - \theta_0, \frac{\pi}{2} + \theta_0)$ is the angle between two instances of M-dimensional unit vectors, and $\theta_0 = arccos\left(\left(\delta\Gamma(\frac{M-1}{2})/\sqrt{\pi}\Gamma(\frac{M}{2})\right)^{1/(M-2)}\right)$. Cai et al. [7] has proved that when $M > 2$, θ concentrates around $\frac{\pi}{2}$, and the concentration becomes stronger as the dimension M grows. In particular, θ converges to $\frac{\pi}{2}$ at the rate of $\sqrt{M}$ when M approaches infinity [7, p. 1840]. Thus, $\frac{1}{sin^2(\theta)}$ has a decreasing trend and converges to 1 due to the symmetry of sine function on $[0, \pi]$. Subsequently, we can have that $C_R = \frac{2}{\varepsilon\psi}$ has a decreasing trend as well. In addition, when M goes large, we get that C_R converges to $\frac{2}{\varepsilon}$ and hence $\mathbb{E}_{R1SMG}[\mathcal{L}]$ can be as low as $\frac{2(\Delta_2 f)^2}{\varepsilon}$[3]. □

Curse Lifted. (15) clearly shows that we have lifted the identified curse by considering noise attributed to a singular multivariate Gaussian distribution with rank-1 covariance matrix. In particular, the expected accuracy loss introduced by the R1SMG mechanism is $Tr[\boldsymbol{\Pi}]$, which equal to the only nonzero eigenvalue σ_* but is **not** the summation of M positive eigenvalues any more like in the existing Gaussian mechanisms.

From Theorem 4 and Theorem 7, we arrive at Corollary 1.

Corollary 1. *The R1SMG mechanism leads to expected accuracy loss on a lower order of magnitude by at least M or MN compared with the classic Gaussian, analytic Gaussian, and MVG mechanisms. In particular, we have $C_R = \Theta(\frac{\varepsilon C_C}{M}) = \Theta(\frac{C_A}{M}) = \Theta(\frac{C_M}{MN})$.*

Proof. According to Theorem 5, we have $\frac{1}{\psi} = \left(\frac{\sqrt{\pi}\Gamma(\frac{M}{2})}{\delta\Gamma(\frac{M-1}{2})}\right)^{\frac{2}{M-2}} = \exp\left(\frac{2}{M-2}\ln\left(\frac{\sqrt{\pi}\Gamma(\frac{M}{2})}{\delta\Gamma(\frac{M-1}{2})}\right)\right)$. By connecting $\frac{\Gamma(\frac{M}{2})}{\Gamma(\frac{M-1}{2})}$ to the normalization constant of the Beta distribution $Beta(\alpha, \beta)$ with $\alpha = \frac{M-1}{2}$ and $\beta = \frac{1}{2}$, it gives $\Theta\left(\frac{\Gamma(\frac{M}{2})}{\Gamma(\frac{M-1}{2})}\right) = \Theta\left(\frac{\sqrt{M-2}}{\sqrt{2}}\right)$ [7]. As a result, we can obtain $C_R = \Theta\left(\frac{2}{\varepsilon}\exp\left(\frac{2}{M-2}\ln\frac{\sqrt{M-2}}{\sqrt{2}}\right)\right)$. Compare C_R with the results in Theorem 4, we get that $C_R = \Theta(\frac{\varepsilon C_C}{M}) = \Theta(\frac{C_A}{M}) = \Theta(\frac{C_M}{MN})$. □

Moreover, in Appendix B, we plot the empirical $\mathbb{E}[\mathcal{L}]$ versus M achieved by the R1SMG mechanism and by other mechanisms, respectively, and clearly show that $\mathbb{E}_{R1SMG}[\mathcal{L}]$ has a decreasing trend as M increases and eventually converges, whereas all the other mechanisms result in expected accuracy loss increasing with M (cf. Figure 9).

5.4 Discussion on Privacy Leakage of Utilizing Vector in the Null Space of v

The R1SMG mechanism does not span the entire space of $\mathbb{R}^M$. Therefore, it may raise concerns about privacy leakage if one takes advantage of the vector in the null space of $\mathbf{v}$. In other words, let $\mathbf{s}$ be the output of R1SMG and $\mathbf{u} \in Null(\mathbf{v}) = \{\mathbf{u} | < \mathbf{u}, \mathbf{v} >= 0\}$. Since $\mathbf{n} \in Span(\mathbf{v})$, one can have $< \mathbf{u}, \mathbf{s} >=< \mathbf{u}, f(\boldsymbol{x}) + \mathbf{n} >=< \mathbf{u}, f(\boldsymbol{x}) >$, which implies potential privacy leakage of queries on $f(\boldsymbol{x})$.

However, we highlight that $\mathbf{v}$ is randomly sampled from $\mathcal{U}(\mathbb{V}_{1,M})$, which is an intermediate noise to produce the final noise of the R1SMG mechanism and will not be made public. Note that this is not against the core idea of DP, i.e., "privacy without obscurity", because the entire noise generation process, i.e, (11), including the choice of distribution parameters are transparent. Thus, the probability that a constructed $\mathbf{u}$ that is in the null space of a randomly sampled $\mathbf{v}$ is zero, i.e., $\Pr[\mathbf{u} \in Null(\mathbf{v})] = \frac{\mu(Null(\mathbf{v}))}{\mu(\mathbb{R}^M)} = 0$, where $\mu(\cdot)$ is the Lebesgue measure of a measurable set.

For better explanation, , we give toy examples in $\mathbb{R}^2$ and $\mathbb{R}^3$ (when $M = 2$ and $M = 3$). Note that in practice R1SMG requires $M > 2$, here we use $\mathbb{R}^2$ just to visualize the idea (the probability that a constructed $\mathbf{u}$ that is in the null space of a randomly sampled $\mathbf{v}$ is zero) which is independent of the dimension. In $\mathbb{R}^2$, $\mathbb{V}_{1,2}$ is the unit circle shown in Figure 3 (left). Suppose that $\mathbf{v}_1 = [-1/\sqrt{5}, 2/\sqrt{5}]^T$ is the randomly sampled variable in $\mathbb{V}_{1,2}$. Then $Null(\mathbf{v}_1)$ is I (the dashed line orthogonal to $\mathbf{v}_1$). Clearly, the probability of sampling a point from $\mathbb{R}^2$ that resides on a specific line is 0. Likewise, in $\mathbb{R}^3$, $\mathbb{V}_{1,3}$ is the unit sphere shown in Figure 3 (right). Suppose

[3] Note that the asymptotic property of C_R is studied when ε is always feasible. Suppose that we are interested in the asymptotic property of C_R when $M_1 < M < M_2$. Then, as explained in Remark 1, a feasible ε refers to $\varepsilon < \frac{1}{M_2}$.

that $\mathbf{v}_2 = [1/\sqrt{3}, -1/\sqrt{3}, 1/\sqrt{3}]^T$ is the randomly sampled variable in $\mathbb{V}_{1,3}$. Then $Null(\mathbf{v}_2)$ is $\mathcal{F}$ (the plane orthogonal to $\mathbf{v}_2$). The probability of sampling a point from $\mathbb{R}^3$ that resides on a specific plane is also 0. Although, it is publicly known that $\mathbf{v}$ is sampled from $\mathcal{U}(\mathbb{V}_{1,M})$, the Lebesgue measure of $Null(\mathbf{v})$ (given $\mathbf{v} \subset \mathbb{V}_{1,M}$) is still 0. For example, the probability of sampling the blue points in Figure 3 (left) is 0, and the probability of sampling points residing on $\mathcal{C} \subset \mathcal{F}$ (the circle orthogonal to $\mathbf{v}_2$) in Figure 3 (right) is also 0. Hence, by generating $\mathbf{v}$ randomly, the R1SMG mechanism will cause privacy leakage of using the vectors in the null space of $\mathbf{v}$ to have probability zero.

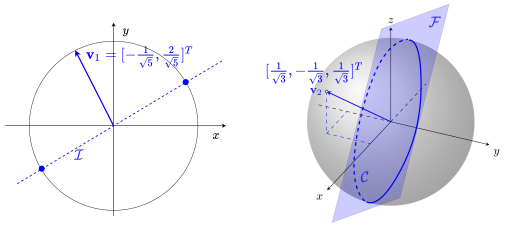

Figure 3: Visualization of the null space of the noise generated by the R1SMG mechanism. Left: $\mathbb{V}_{1,2}$. Right: $\mathbb{V}_{1,3}$.

6 Accuracy Stability

Now, we evaluate the accuracy stability achieved by various output perturbation DP mechanisms by studying the kurtosis and skewness of the distribution of $\mathcal{L}$ (the non-deterministic accuracy loss defined in Definition 1). In particular,

- Kurtosis, a descriptor of "tail extremity" of a probability distribution, is defined as the ratio between the 4th moment and the square of the 2nd moment of a random variable, i.e., $\frac{\mathbb{E}[\mathcal{L}^4]}{(\mathbb{E}[\mathcal{L}^2])^2}$. A larger kurtosis means that outliers or extreme large values are less likely to be generated in a given probability distribution [45].
- Skewness, a descriptor of the "bulk" of a probability distribution, is defined as the ratio between the 3rd moment and the square root of the cube of the 2nd moment of a random variable, i.e., $\frac{\mathbb{E}[\mathcal{L}^3]}{(\mathbb{E}[\mathcal{L}^2])^{3/2}}$. A larger skewness means that the bulk of the samples is at the left region of the PDF and the right tail is longer.

As a result, in order to have high accuracy stability on the perturbed query result, $\mathcal{L}$ with both **larger kurtosis and skewness is preferred**. We summarize the theoretical results for various mechanisms in Theorem 8 and Theorem 9.

Theorem 8. *The kurtosis of the distribution of $\mathcal{L}$ in the R1SMG mechanism is $\frac{35}{3}$ which is larger than that of the classic Gaussian mechanism, of the analytic Gaussian mechanism, and of the MVG mechanism, i.e., the PDF of $\mathcal{L}$ in the R1SMG mechanism is more leptokurtic than that in the classic Gaussian, in the analytic Gaussian, and in the MVG mechanism.*

Proof. First, for the proposed R1SMG, we have

$$\text{Kurt}_{\text{R1SMG}}(\mathcal{L}) = \frac{\mathbb{E}[\mathcal{L}^4]}{(\mathbb{E}[\mathcal{L}^2])^2} \overset{(a)}{=} \frac{\mathbb{E}\left[\left((\mathbf{v}\sqrt{\sigma_*}z)^T\mathbf{v}\sqrt{\sigma_*}z\right)^4\right]}{\left(\mathbb{E}\left[\left((\mathbf{v}\sqrt{\sigma_*}z)^T\mathbf{v}\sqrt{\sigma_*}z\right)^2\right]\right)^2} = \frac{\mathbb{E}[(z\sigma_* z)^4]}{\left(\mathbb{E}[(z\sigma_* z)^2]\right)^2}$$
$$\overset{(b)}{=} \frac{48\sigma_*^4+32\sigma_*^3\sigma_*+12(\sigma_*^2)^2+12\sigma_*^2(\sigma_*)^2+(\sigma_*)^4}{\left(2\sigma_*^2+(\sigma_*)^2\right)^2} = \frac{35}{3},$$

where (a) is due to the noise generation process of R1SMG in (11) and (b) is obtained by applying Lemma 1 with $\mathbf{A} = 1$ and $\mathbf{\Xi} = \sigma_*$ (i.e., the case of univariant Gaussian).

Then, we show that both classic Gaussian mechanism and the MVG mechanism will result in $\mathcal{L}$ with kurtosis less than $\frac{35}{3}$. In particular, the classic Gaussian mechanism (denoted as $\mathcal{G}$) adds i.i.d. noise to each entry of $f(\boldsymbol{x})$, $\mathbf{n}_i \sim \mathcal{N}(0, \sigma^2), i \in [1, M]$. Thus, we have $\mathbb{E}_{\mathcal{G}}[\mathcal{L}^t] = \mathbb{E}_{z_i \sim \mathcal{N}(0,1)}\left[\left(\sum_i \sigma^2 z_i^2\right)^t\right] = \sigma^{2t}\mathbb{E}_{U\sim\chi^2(M)}[U^t] = \sigma^{2t}2^t\frac{\Gamma(t+\frac{M}{2})}{\Gamma(\frac{M}{2})}$, where the last equality follows from the tth moments of Chi-squared random variable. Set t to be 4 and 2, one can verify that $\text{Kurt}_{\mathcal{G}}(\mathcal{L}) = \frac{\mathbb{E}_{\mathcal{G}}[\mathcal{L}^4]}{(\mathbb{E}_{\mathcal{G}}[\mathcal{L}^2])^2} = \sigma^8 2^4\frac{\Gamma(4+\frac{M}{2})}{\Gamma(\frac{M}{2})} \Big/ \left(\sigma^4 2^2\frac{\Gamma(2+\frac{M}{2})}{\Gamma(\frac{M}{2})}\right)^2 < \frac{35}{3}, \forall M > 1$.

The same analysis holds for the analytic Gaussian mechanism, as it adds i.i.d. Gaussian noise with variance σ_A^2.

MVG introduces the matrix-valued noise attributed to a matrix-valued Gaussian distribution, i.e., $\mathcal{N}(\mathbf{0}, \mathbf{\Sigma}, \mathbf{\Psi})$, and the vectorization of the noise matrix is also attributed to a multivariate Gaussian distribution, i.e., $\mathcal{N}(\mathbf{0}, \mathbf{\Sigma}\otimes\mathbf{\Psi})$ [21]. Denote $\mathbf{\Xi} = \mathbf{\Sigma}\otimes\mathbf{\Psi}$, and let $k > 1$ be an integer, we first have

$$(Tr[\mathbf{\Xi}])^k = (Tr[\mathbf{\Sigma}\otimes\mathbf{\Psi}])^k = \left[\left(\sum_{i=1}^{\text{rank}(\mathbf{\Sigma})}\sigma_i(\mathbf{\Sigma})\right)\left(\sum_{j=1}^{\text{rank}(\mathbf{\Psi})}\sigma_j(\mathbf{\Psi})\right)\right]^k$$
$$> \sum_{i=1,j=1}^{\text{rank}(\mathbf{\Sigma}),\text{rank}(\mathbf{\Psi})}\left(\sigma_i(\mathbf{\Sigma})\sigma_j(\mathbf{\Psi})\right)^k = Tr\left[(\mathbf{\Sigma}\otimes\mathbf{\Psi})^k\right] = Tr\left[\mathbf{\Xi}^k\right].$$

Similarly, for any positive integer $k_1 > k_2 \geq 1$ and symmetric $\mathbf{\Xi} \in \mathbb{PD}$, we have

$$Tr[\mathbf{\Xi}^{k_1}](Tr[\mathbf{\Xi}])^{k_2} > Tr[\mathbf{\Xi}^{k_1}]Tr[\mathbf{\Xi}^{k_2}]$$
$$= \left(\sum_{i=1}^{\text{rank}(\mathbf{\Xi})}(\sigma_i(\mathbf{\Xi}))^{k_1}\right)\left(\sum_{i=1}^{\text{rank}(\mathbf{\Xi})}(\sigma_i(\mathbf{\Xi}))^{k_2}\right) > \sum_{i=1}^{\text{rank}(\mathbf{\Xi})}(\sigma_i(\mathbf{\Xi}))^{k_1+k_2}$$
$$= Tr[\mathbf{\Xi}^{k_1+k_2}].$$

Then, according to Lemma 1, we have

$$\text{Kurt}_{\text{MVG}}(\mathcal{L}) = \frac{\mathbb{E}[\mathcal{L}^4]}{(\mathbb{E}[\mathcal{L}^2])^2} = \frac{\mathbb{E}[(\mathbf{n}^T\mathbf{n})^4]}{(\mathbb{E}[(\mathbf{n}^T\mathbf{n})^2])^2} = \frac{\mathbb{E}[Q(\mathbf{n})^4]}{(\mathbb{E}[Q(\mathbf{n})^2])^2}$$
$$= \frac{48Tr[\mathbf{\Xi}^4]+32Tr[\mathbf{\Xi}^3]Tr[\mathbf{\Xi}]+12(Tr[\mathbf{\Xi}^2])^2+12Tr[\mathbf{\Xi}^2](Tr[\mathbf{\Xi}])^2+(Tr[\mathbf{\Xi}])^4}{\left(2Tr[\mathbf{\Xi}^2]+(Tr[\mathbf{\Xi}])^2\right)^2}.$$

Next, we can obtain $48Tr[\mathbf{\Xi}^4] < \frac{104}{3}(Tr[\mathbf{\Xi}^2])^2 + \frac{8}{3}Tr[\mathbf{\Xi}^2](Tr[\mathbf{\Xi}])^2 + \frac{32}{3}(Tr[\mathbf{\Xi}])^4$ and $32Tr[\mathbf{\Xi}^3]Tr[\mathbf{\Xi}] < 32Tr[\mathbf{\Xi}^2](Tr[\mathbf{\Xi}])^2$. Thus, we get $\text{Kurt}_{\text{MVG}}(\mathcal{L}) < \frac{\frac{35}{3}\left(4Tr[\mathbf{\Xi}^2]+4Tr[\mathbf{\Xi}^2](Tr[\mathbf{\Xi}])^2+(Tr[\mathbf{\Xi}])^4\right)}{4Tr[\mathbf{\Xi}^2]+4Tr[\mathbf{\Xi}^2](Tr[\mathbf{\Xi}])^2+(Tr[\mathbf{\Xi}])^4} = \frac{35}{3}$, which concludes the proof. □

Theorem 9. *The skewness of the distribution of $\mathcal{L}$ in the R1SMG mechanism is $\frac{5\sqrt{3}}{3}$ which is larger than that of the classic Gaussian mechanism, of the analytic Gaussian mechanism, and of the MVG mechanism, i.e., the PDF of $\mathcal{L}$ in the R1SMG mechanism is more right-skewed than that in the classic Gaussian, in the analytic Gaussian, and in the MVG mechanism.*

Proof. The proof follows the same procedure as the proof of Theorem 8, thus we only show the key steps here.

For the proposed R1SMG, we have $\text{skew}_{\text{R1SMG}}(\mathcal{L}) = \frac{\mathbb{E}[\mathcal{L}^3]}{(\mathbb{E}[\mathcal{L}^2])^{3/2}} = \frac{\mathbb{E}\left[(z\sigma_* z)^3\right]}{\left(\mathbb{E}\left[(z\sigma_* z)^2\right]\right)^{3/2}} = \frac{\sigma_*^3 2^3 \frac{\Gamma(3+\frac{1}{2})}{\Gamma(\frac{1}{2})}}{\left(\sigma_*^2 2^2 \frac{\Gamma(2+\frac{1}{2})}{\Gamma(\frac{1}{2})}\right)^{3/2}} = \frac{5\sqrt{3}}{3}$.

For the classic Gaussian mechanism, one can verify that $\text{skew}_{\mathcal{G}}(\mathcal{L}) = \frac{\sigma_G^6 2^3 \frac{\Gamma(3+\frac{M}{2})}{\Gamma(\frac{M}{2})}}{\left(\sigma_G^4 2^2 \frac{\Gamma(2+\frac{M}{2})}{\Gamma(\frac{M}{2})}\right)^{3/2}} = \frac{\frac{\Gamma(3+\frac{M}{2})}{\Gamma(\frac{M}{2})}}{\left(\frac{\Gamma(2+\frac{M}{2})}{\Gamma(\frac{M}{2})}\right)^{3/2}} < \frac{5\sqrt{3}}{3}, \forall M > 1$.

The same analysis holds for the analytic Gaussian mechanism, as it adds i.i.d. Gaussian noise with variance σ_A^2.

For the MVG mechanism, it is easy to check that $\text{skew}_{\text{MVG}}(\mathcal{L}) = \frac{\mathbb{E}[\mathcal{L}^3]}{(\mathbb{E}[\mathcal{L}^2])^{3/2}} = \frac{\mathbb{E}[(\mathbf{n}^T\mathbf{n})^3]}{(\mathbb{E}[(\mathbf{n}^T\mathbf{n})^2])^{3/2}} = \frac{\left(8Tr[\Xi^3]+6Tr[\Xi^2]Tr[\Xi]+(Tr[\Xi])^3\right)}{\left(2Tr[\Xi^2]+(Tr[\Xi])^2\right)^{3/2}}$, $\Xi = \mathbf{\Sigma} \otimes \mathbf{\Psi}$. Besides, we have shown that for any positive integer k, k_1, and k_2, we have $(Tr[\Xi])^k > Tr[\Xi^k]$ and $Tr[\Xi^{k_1}](Tr[\Xi])^{k_2} > Tr[\Xi^{k_1+k_2}]$. Thus, one can check that $(\text{skew}_{\text{MVG}}(\mathcal{L}))^2 = \frac{\left(8Tr[\Xi^3]+6Tr[\Xi^2]Tr[\Xi]+(Tr[\Xi])^3\right)^2}{\left(2Tr[\Xi^2]+(Tr[\Xi])^2\right)^3} < \frac{\frac{25}{3}\left(2Tr[\Xi^2]+(Tr[\Xi])^2\right)^3}{\left(2Tr[\Xi^2]+(Tr[\Xi])^2\right)^3}$, i.e., $\text{skew}_{\text{MVG}}(\mathcal{L}) < \frac{5\sqrt{3}}{3}$. □

Remark 2. *We can consider the univariate Gaussian with unit covariance, i.e., $\mathcal{N}(\mu, 1)$, as a reference distribution, whose kurtosis value and skewness value are 3 and 0, respectively. Then, it suggests that the distribution of $\mathcal{L}$ in the R1SMG mechanism are much more leptokurtic and right-skewed than $\mathcal{N}(\mu, 1)$. Moreover, since the noise used in the classic Gaussian, analytic Gaussian, and MVG mechanisms are characterized by full-rank covariance matrices, their kurtosis values and skewness values asymptotically converge to $3 + 12/H$ and $\sqrt{8/H}$, respectively [25] (H is the degree of freedom of the obtained $\mathcal{L}$, and $H = M$ for $f(\boldsymbol{x}) \in \mathbb{R}^M$). It means that the distributions of $\mathcal{L}$ in the Gaussian mechanism and MVG are similar to $\mathcal{N}(\mu, 1)$ when M is large.*

In Section 7.1, by using 2D count query as an example, we will empirically show that the distribution of $\mathcal{L}$ obtained in the R1SMG mechanism is more leptokurtic and right-skewed than those of the classic Gaussian, the analytic Gaussian, and the MVG mechanisms (cf. Figure 5). In other words, the R1SMG mechanism can provide differentially private query results with the highest accuracy stability.

Based on the above theoretical analysis, we can have the following corollary.

Corollary 2. *The R1SMG mechanism outperforms the classic Gaussian, the analytic Gaussian, and the MVG mechanism, because it boosts the utility of the query results by achieving lower expected accuracy loss and higher accuracy stability.*

7 Experiments

In this section, we conduct three case studies to validate the utility boosting achieved by the R1SMG mechanism, i.e., 2D count query, principal component analysis (PCA), and deep learning, all in a differentially private manner. The query results of these case studies are either matrices or tensors, so the R1SMG, classic Gaussian, and analytic Gaussian mechanisms will first generate noise in vector form, and then reshape the noise into the forms required by different studies.

7.1 Case Study I: Uber Pickup Count Query

In this case study, we utilize the New York City (NYC) Uber pickups dataset [2], on which we are interested in releasing the counts of Uber pickups in different areas of NYC from "4/1/2014 00:11:00" to "4/3/2014 23:57:00" in a differentially private manner. The size of the count query $f(\boldsymbol{x})$ is determined by the partition areas (defined later) of NYC. Location and trajectory privacy breaches have been reported and investigated in several research works [23, 32, 35, 46]. Thus, it is important to guarantee that the existence or absence of any pickup record is private when sharing the count query.

In particular, we partition the map of NYC into small areas by evenly dividing the latitude and longitude into 89 disjoint intervals. As a result, the considered query is $f(\boldsymbol{x}) \in \mathbb{R}^{89\times 89}$, whose entries record the number of Uber pickups in different small areas. The squared sensitivity of this query is 2, because the absence or presence of a specific pickup can change at most 2 entries of $f(\boldsymbol{x})$ by 1. We also consider that the data consumer (i.e., the query result recipient) has the prior knowledge of the valid pickup areas in the resulted 89×89 small areas. Thus, he can perform post-processing on the noisy count query to eliminate the values in invalid areas (e.g., it is impossible to have Uber pickups over the Hudson River). We visualize the non-private Uber pickup count query in Figure 4(a). We observe high volumes of pickups around Soho, Fifth avenue, and LaGuardia Airport.

Comparisons with Mechanisms. In additional to the (ε, δ)-DP output perturbation mechanisms, i.e., the classic Gaussian, analytic Gaussian, MVG, and MGM (a variant of MVG [47]) mechanisms, we also compare the R1SMG mechanism with the mechanisms that are specially developed for differentially private 2D count queries. In particular, they are (i) DAWA [31], which obtains differentially private count queries by adding noise that is adapted to both the input data

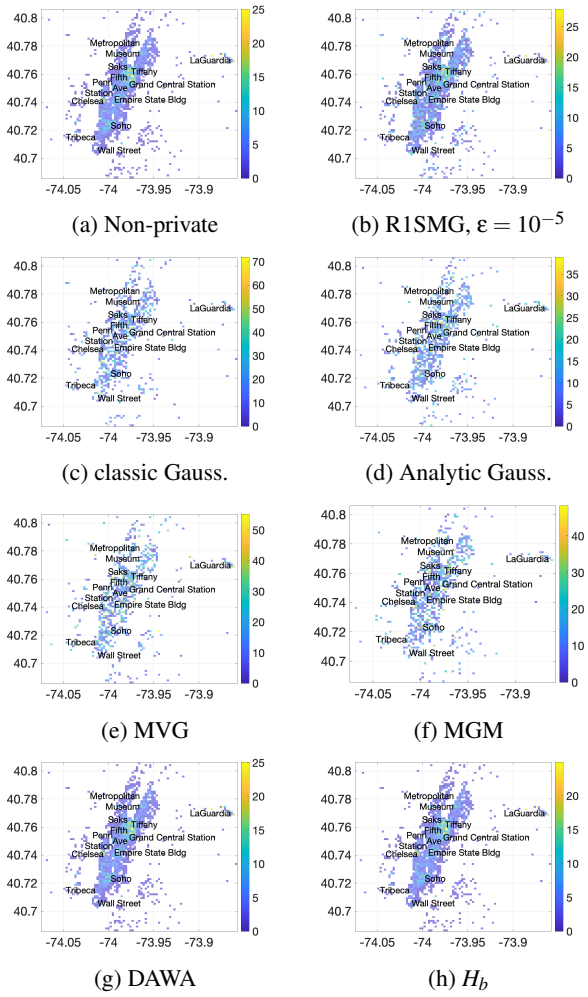

(a) Non-private (b) R1SMG, $\varepsilon = 10^{-5}$

(c) classic Gauss. (d) Analytic Gauss.

(e) MVG (f) MGM

(g) DAWA (h) H_b

Figure 4: Visualization of (a) non-private counts, (b)-(h) are differentially private 2D counts obtained by the R1SMG, classic Gaussian, analytic Gaussian, MVG, MGM, DAWA, and H_b mechanisms, respectively. ε is 10^{-5} for the R1SMG mechanism and is 0.5 for the other mechanisms.

and the given query set, and (ii) H_b [38], which answers range queries using noisy hierarchies of equi-width histograms.

Results. We visualize the differential-privately released count queries obtained by the R1SMG and the other comparing mechanisms in Figures 4(b)-(h). The privacy budgets for all the comparing mechanisms are $\varepsilon = 0.5$ and $\delta = 10^{-7}$. For the R1SMG mechanism, we set $\varepsilon = 10^{-5} < \frac{0.5}{89^2}$ according to Remark 1.

Clearly, the R1SMG mechanism outperforms all the comparing output perturbation mechanisms (i.e., Figure 4(c)-(f)) even under a very restricted privacy budget, as it preserves the global and local patterns of the non-private count query in Figure 4(a) in the best way. In contrast, the classic Gaussian, MVG, and MGM mechanisms greatly compromise the utility of the count query results although their privacy budget is 10^4 times as large as that of our R1SMG mechanism. For example, in Figures 4(c), (e), and (f) many areas on the Manhattan island have no pickup counts, and the maximum count is also increased significantly (i.e., higher than 60, 50, and 40, respectively, compared to the non-private maximum count which is only around 27). Although the analytic Gaussian mechanism outperforms the classic Gaussian and MVG mechanisms by calibrating the noise variance exactly, it still introduces noise with higher magnitude than the R1SMG mechanism. Finally, we observe that the R1SMG mechanism has comparable performance with DAWA and H_b that are developed specially for differentially private range queries. Note that DAWA has slightly better results because it is data-dependent (the noise is customized to the input data [31]), and that H_b adopts "constrained inference" (an accuracy boosting scheme) to control the mean square error of counting queries at various granularities. Moreover, not only does the R1SMG mechanism introduces comparable errors with those by the two task-specific mechanisms under a much stricter privacy guarantee, but also it is very promising in other real-world tasks because it is not task-dependent.

Utility boosting and stability. Next, we empirically corroborate that the R1SMG mechanism boosts the utility of the NYC Uber pickup count query by achieving lower accuracy loss and higher accuracy stability. Specifically, by setting $\varepsilon = 10^{-5}$ for the R1SMG mechanism and $\varepsilon = 0.5$ for the other mechanisms and $\delta = 10^{-7}$, we repeat the noise generation process 100 times for the R1SMG, classic Gaussian, analytic Gaussian, and MVG mechanisms, respectively, and calculate the accuracy loss (i.e., $\mathcal{L}$, in Definition 1). In Figure 5, we plot the corresponding histograms, empirical average values of $\mathcal{L}$, and the empirical PDFs, obtained by different mechanisms. In particular, Figure 5(a) plots the histograms of $\mathcal{L}$ for various mechanisms on the same x-axis, and Figure 5(b) is the zoomed in histograms for each mechanism. Figure 5(c) demonstrates the empirical PDF and mean of $\mathcal{L}$ for each mechanism on the same x-axis, and Figure 5(d) shows the corresponding zoomed in plots.

Figure 5 shows that the empirical mean accuracy loss obtained by the R1SMG mechanism is less than those obtained by the other mechanisms, in spite of the fact that the privacy budget is smaller than $1/10^4$ of the others. This means that under very restricted privacy guarantee, the R1SMG mechanism is able to introduce additive noise with small magnitude (squared Frobenius norm), which lead to lower empirical accuracy loss. In particular, $\mathcal{L}$ obtained by the R1SMG mechanism is around 1.5×10^4 when $\varepsilon = 10^{-5}$, whereas, $\mathcal{L}$ obtained by the other mechanisms are around 5×10^4 even with $\varepsilon = 0.5$. We can also find that the histogram and the empirical PDF of $\mathcal{L}$ obtained by the R1SMG mechanism are much more leptokurtic and right-skewed than those achieved by the others. Particularly, in Figure 5(d), compared with those obtained by the other mechanisms that have bell shapes, the empirical PDF of $\mathcal{L}$ achieved by the R1SMG mechanism has a longer and fatter tail, a higher and sharper central peak, and the den-

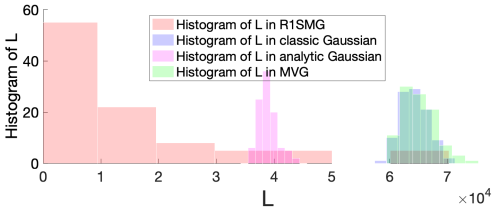

(a) Histogram of $\mathcal{L}$ for various mechanisms

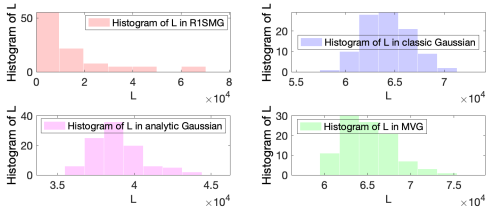

(b) Histogram of $\mathcal{L}$ (zoomed in)

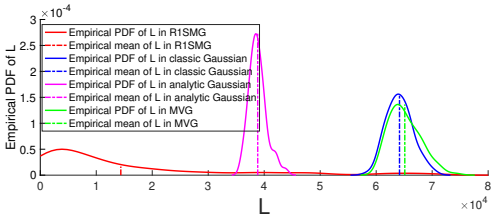

(c) Empirical PDF and mean of $\mathcal{L}$ for various mechanisms

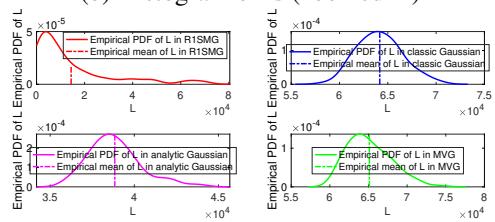

(d) Empirical PDF and mean of $\mathcal{L}$ (zoomed in)

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

sity is more densely distributed on the left side. The results corroborate the analysis in Remark 2. Thus, under very strict (ε, δ)-DP requirement, it is still less likely for the R1SMG mechanism to generate additive noise with large magnitude.

7.2 Case Study II: PCA

In this case study, we explore differentially private principal component analysis (PCA) using the Swarm Behaviour dataset [1]. It contains 24016 records of flocking behaviour (i.e., the way that groups of birds, insects, fish or other animals, move close to each other). Each record has 2400 features characterizing a flocking observation, e.g., radius of separation of groups of insects, moving direction, moving velocity, etc. Each data record is assigned a binary label, and "1" represents "flocking behavior" and "0" represents "not flocking". The values of all features are normalized between -1 and 1. Similar to [11], we consider the query $f(\boldsymbol{x})$ as the covariance matrix of this dataset, i.e., $f(\boldsymbol{x}) = D^T D/24016 \in \mathbb{R}^{2400\times 2400}$, where $D \in \mathbb{R}^{24016\times 2400}$ denotes the considered Swarm Behaviour data.[4] Then, the l_2 sensitivity of the query function is $\Delta_2 f = \sup_{D_1,D_2} \frac{||D_1(i)^T D_1(i) - D_2(i)^T D_2(i)||_F}{24016} = \frac{\sqrt{2\times 2400^2}}{24016} = \frac{2400\sqrt{2}}{24016}$.

We conduct singular value decomposition on the queried result (i.e., the noisy empirical covariance matrix) to extract the components of the dataset. The considered evaluation metric is the total deviation which is defined as $\Delta = \sum_{i=1}^{2041} \Delta_i = \sum_{i=1}^{2041} |\lambda_i - \widetilde{\mathbf{v}}_i^T \mathbf{C} \widetilde{\mathbf{v}}_i|$, where $\mathbf{C}$ is the ground-truth (non-private) empirical covariance matrix, λ_i is the ith eigenvalue of $\mathbf{C}$, and $\widetilde{\mathbf{v}}_i$ is the ith differentially private component (eigenvector). Specifically, Δ_i quantifies the deviation of the variance from λ_i in the direction of the ith component. The smaller Δ is, the higher utility the obtained components have, and we have $\Delta = 0$ for the non-private baseline.

Since PCA usually serves as a precursor to classification task, we also evaluate the utility of the queried data by projecting the dataset onto the subspace spanned by the first 20 PCs obtained from various noisy covariance matrices, and then measure the classification accuracy using the projected data. The higher the classification accuracy, the higher the utility of the perturbed covariance matrix. In the experiments, we use 60% and 40% of the dataset for training and testing, respectively, and the classification algorithm is the linear SVM. Note that the accuracy of the non-private baseline (i.e., classification using the original dataset project onto the first 20 PCs of the original covariance matrix) is 96.89%.

Comparisons with other Mechanisms. In addition to the output perturbation mechanisms, we also consider the principal components obtained from 3 task-specific differentially private algorithms designed specially for PCA, i.e., (i) PPCA [12], which generates privacy-preserving components by sampling orthonormal matrices from the matrix Bingham distribution parametermized by the original non-private empirical covariance matrix scaled by the privacy budget, (ii) MOD-SULQ [12], which directly perturbs the covariance matrix using calibrated Gaussian noise, and (iii) the Wishart mechanism [24], which perturbs the empirical covariance matrices using noise attributed to the Wishart distribution.

Results. In this experiment, we vary ε from 0.2 to 1.6 and set $\delta = 10^{-7}$ for all the comparing mechanisms and let the privacy budget of the R1SMG be only $1/10^7$ of the others according to the discussion in Remark 1.

Figure 6(a) and (b) plot the total derivation (Δ) and classification accuracy on the projected testing data obtained from all mechanisms. Obviously, the R1SMG mechanism has high utility although its the privacy budget is extremely limited, e.g.,

[4] Since covariance matrices are symmetric, the data consumer will perform post processing $\left(\mathcal{M}(f(\boldsymbol{x})) + \mathcal{M}(f(\boldsymbol{x}))^T\right)/2$ to obtain symmetric result. The privacy guarantee naturally holds due to the post-processing immunity [17].

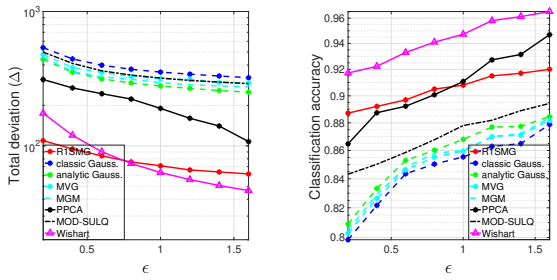

(a) Deviation versus ε. (b) Accuracy versus ε.

Figure 6: Utility of the noisy covariance matrices obtained by various mechanisms. (a) The total deviation. (b) The classification accuracy measured on the projected testing dataset. ε used in the R1SMG mechanism is only $1/10^7$ of the others.

$\varepsilon < 10^{-7}$, which is also close to the non-private baselines (i.e., $\Delta = 0$ and 96.89% accuracy). This is because the R1SMG mechanism introduces perturbation noise of much smaller magnitudes in spite of very strict privacy guarantee, compared with the other mechanisms. Taking the MVG mechanism for example, according to Theorem 3, even when $\varepsilon = 1.6$ and $\delta = 10^{-7}$, MVG requires $||\sigma(\mathbf{\Sigma}^{-1})||_2||\sigma(\mathbf{\Psi}^{-1})||_2 \leq 0.011$ to achieve the desired privacy guarantee. This means that the covariance matrices will have very large singular values, which implies large expected accuracy loss (as discussed in Section 3) and causes the original data be overwhelmed by the additive noise. In fact, according to [11], for the MVG mechanism to have a decent performance on differentially private PCA, it requires $\varepsilon \geq 5$. Moreover, in this case study, the R1SMG mechanism even has similar utility performance of that of the task-specific mechanisms including PPCA, MOD-SULQ, and Wishart. In other words, under very limited ε, the R1SMG mechanism not only achieves very small Δ, but also gives high classification accuracy. This suggests that R1SMG is very promising in boosting the accuracy of high-dimensional differentially private query even if $\varepsilon \ll 1$. Note that in this study, the R1SMG mechanism also achieves the most stable accuracy, i.e., the accuracy loss introduced by it has leptokurtic and right-skewed histogram and empirical PDF, but the others still have bell-shaped empirical PDFs (similar to Figure 5). We omit the plots due to space limit.

7.3 Case Study III: Deep Learning with DP

In this case study, we conduct preliminary validation to assess the feasibility of replacing the classic Gaussian mechanism with the R1SMG in the DPSGD (differentially private stochastic gradient descent) optimization framework [3], and demonstrate that R1SMG can boost the utility (testing accuracy) of a differentially-private deep learning model and improve its training efficiency by reducing the number of epochs. Please note that this case study is not intended to be thorough or comprehensive, as differentially-private deep learning constitutes a distinct research direction in its own right.

As a proof of concept, we consider the MNIST [30] and CIFAR-10 [28] dataset, and compare our R1SMG-based DPSGD with the seminal work of DPSGD in deep learning (Abadi et al. CCS'16 [3]). For a given deep neural network, Abadi et al. introduce calibrated i.i.d. Gaussian noise ($\mathcal{N}(\mathbf{0}, \sigma^2 C^2 \mathbf{I})$) to the gradients computed on grouped batches of training data, where C is a norm clipping parameter to control the sensitivity (Algorithm 1 in [3]). They also developed the privacy accountant scheme to tightly bound the cumulative privacy loss for the entire training process. In particular, the total privacy budget (ε, δ) is determined by noise level σ, sampling ratio q, and the number of epochs E (so number of gradient descent steps is $T = E/q$).

In the experiment, we give our R1SMG-based DPSGD the same total privacy budget (ε, δ), C, q, and E as those of DPSGA in [3]. Then, we amortize (ε, δ) to all gradient descent steps in our R1SMG-based DPSGD, i.e., each step is assigned with $(\varepsilon_0, \delta_0)$. Although ε_0 might go beyond the upper bound discussed in Remark 1, it makes the comparison with DPSGD clear and straightforward. To be more specific, $(\varepsilon_0, \delta_0)$ is obtained by solving the following equations formulated using the strong composition theorem [19] considering $\frac{E}{q}$ repetitive executions of gradient descent.

$$\begin{cases} \varepsilon = \sqrt{2\frac{E}{q}\ln\left(\frac{1}{\delta'}\right)}(q\varepsilon_0) + \frac{E}{q}(q\varepsilon_0)\left(e^{(q\varepsilon_0)} - 1\right) \\ \delta = \frac{E}{q}(q\delta_0) + \delta' \end{cases}, \quad (16)$$

where $q\varepsilon_0$ and $q\delta_0$ are due to privacy amplification via sampling [6, 26] (see discussion on page 3, right column of [3]). For example, based on the TensorFlow tutorial for DPSGD on MNIST [36], when $C = 1$, $\sigma = 1.1$, $q = \frac{256}{60000}$, and $E = 30$, it leads to $\varepsilon = 1.795$. Then, setting $\delta = 10^{-5}$, $\delta' = 10^{-10}$ and solving (16), we have $(\varepsilon_0, \delta_0) = (0.71, 3.33 \times 10^{-7})$. When $C = 1$, $\sigma = 0.5$, $q = \frac{256}{60000}$, and $E = 30$, it leads to $(16.983, 10^{-5})$-DP for DPSGD, which makes $(\varepsilon_0, \delta_0) = (5.42, 3.33 \times 10^{-7})$ for each gradient descent step in our R1SMG-based DPSGD.

Note that using strong composition theorem to assign amortized privacy budget to each step of our R1SMG-based DPSGD gives more favor to the classic Gaussian mechanism based DPSGD in the comparison. This is because strong composition is much looser than the privacy accountant scheme used in DPSGD (empirically validated by Figure 2 of [3]). Thus, the amortized $(\varepsilon_0, \delta_0)$ can be limited. Developing a tight composition framework for the R1SMG mechanism in SGD-based learning will be a separate research topic.

Even though the original DPSGD is given more favor in the comparison, our R1SMG-based DPSGD still outperforms it by achieving higher training and testing accuracy. We first show the comparison results on MNIST using the above parameters setups in Figure 7. In particular, Figure 7(a) and (b) show the training and testing accuracy on MNIST for the first 30 epochs achieved by various approaches when the entire learning process is $(1.795, 10^{-5})$-DP and $(16.983, 10^{-5})$-DP

(parameters discussed above), respectively. The blue lines and dashed lines in Figure 7 are the result of non-private SGD (the one uses original gradients calculated from the data).

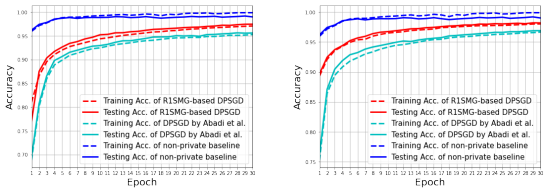

(a) The entire process is $(1.795, 10^{-5})$-DP.

(b) The entire process is $(16.983, 10^{-5})$-DP.

Figure 7: Training and testing acc. on MNIST.

Compared with DPSGD, R1SMG-based DPSGD achieves higher training and testing accuracy in just a few epochs for the MNIST dataset. For example, when the total privacy budget is $(16.983, 10^{-5})$, our method achieves 95% accuracy on the testing dataset using 5 epochs, whereas DPSGD requires 14 epochs. Our R1SMG-based DPSGD also leads to performance that is closer to the non-private baseline, because even though the amortized privacy budget for each R1SMG-based DPSGD step is limited, the magnitude of the perturbation noise introduced by the R1SMG mechanism is much less than that required by the classic Gaussian in DPSGD.

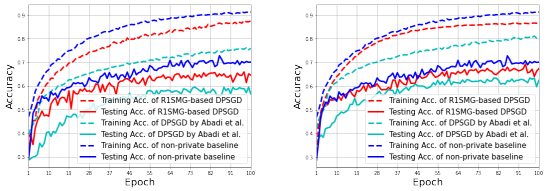

(a) The entire process is $(2.673, 10^{-5})$-DP.

(b) The entire process is $(20.713, 10^{-5})$-DP.

Figure 8: Training and testing acc. on CIFAR-10.

Then, we conduct the experiments on CIFAR-10 dataset with the same parameter setups while using 100 epochs, i.e., $E = 100$. We also adopt the network architecture from the TensorFlow CNN tutorial [44]. The privacy budget for each step of R1SMG-based DPSGD, i.e., $(\varepsilon_0, \delta_0)$ is also decided by solving (16) numerically. In particular, when $\sigma = 1.1$ and $\sigma = 0.5$, the entire process of DPSGD is $(2.673, 10^{-5})$-DP and $(20.713, 10^{-5})$-DP, respectively, which make $(\varepsilon_0, \delta_0)$ become $(0.54, 10^{-7})$ and $(3.52, 10^{-7})$, respectively, for R1SMG-based DPSGD. The comparison results are shown in Figure 8. Clearly, our R1SMG-based DPSGD still outperforms DPSGD, i.e., both training and testing accuracy achieved by R1SMG-based DPSGD are closer to the non-private baselines compared with those achieved by DPSGD. In particular, R1SMG-based DPSGD improve the training efficiency by achieving a higher training accuracy (e.g., $> 80\%$) in much fewer epochs compared with the original DPSGD.

8 Related Work

Many works have attempted to improve the classic Gaussian mechanism. Here, we review some representative ones.

The analytic Gaussian mechanism proposed by Balle et al. [5] (see Theorem 2) improves the classic Gaussian mechanism by calibrating the variance of the Gaussian noise directly using the Gaussian cumulative density function instead of the tail bound approximation, and develop an analytical solution to the variance given specific choices ε and δ. Although this mechanism can reduce the magnitude of the noise when perturbing a scalar-valued quantity, it still suffers the curse of full-rank covariance matrices when applied to perturb high-dimensional query result (as discussed in Section 3).

Zhao et al. [48] investigate the classic Gaussian mechanism under high privacy budgets, i.e., $\varepsilon > 1$, and derive a closed-form upper bounds for the noise variance used in the analytic Gaussian mechanism. Their mechanism can achieve utility higher than the classic Gaussian mechanisms, but still lower than the analytic Gaussian mechanism.

Chanyaswad et al. [11] (see Definition 4 and Theorem 3) propose the MVG mechanism and develop the technique of directional noise to restrict the impact of the perturbation noise on the utility of a matrix-valued query function. However, to obtain the directional noise, one either needs a domain expert to determine the principal components of the queried data or has to calculate them via principal component analysis which consumes some privacy budget.

Some other variants of the classic Gaussian mechanism includes the Generalized Gaussian mechanism [33], which investigates noise calibration under the general l_p sensitivity, and the discrete Gaussian mechanism [9], which uses noise attributed to discrete Gaussian distribution to protect the privacy of query results in a discrete domain.

Our proposed R1SMG mechanism differs with the above mechanisms, as we explicitly require the covariance matrix of the additive noise to be rank-1. Thus, the R1SMG mechanism lifts the identified curse and boosts the utility of query results leading to guaranteed DP at much lower cost of accuracy loss.

9 Conclusions

In this paper, we developed a novel DP mechanism, i.e., R1SMG, to boost the utility and stability of differentially-private query results. We first identify the curse of full-rank covariance matrices in all existing Gaussian-based DP mechanisms. Then, we lift the curse by developing R1SMG which perturbs the query results using noise that follows a singular multivariate Gaussian distribution with a random rank-1 covariance matrix. We rigorously analyze the privacy guarantee of the R1SMG mechanism, and theoretically demonstrate that it can achieve much lower accuracy loss, particularly, on a lower order of magnitude by at least M or MN, and much higher accuracy stability as well, than the classic Gaussian, analytic Gaussian, and the MVG mechanisms.

Acknowledgement

This work was supported in part by the National Science Foundation under Grants EEC-2133630 and CNS-2125460.

References

[1] Swarm behaviour dataset. https://archive.ics.uci.edu/ml/datasets/Swarm+Behaviour.

[2] Uber pickups in new york city. https://www.kaggle.com/fivethirtyeight/uber-pickups-in-new-york-city.

[3] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC conference on computer and communications security, pages 308–318, 2016.

[4] Mir M Ali. Characterization of the normal distribution among the continuous symmetric spherical class. Journal of the Royal Statistical Society: Series B (Methodological), 42(2):162–164, 1980.

[5] Borja Balle and Yu-Xiang Wang. Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In International Conference on Machine Learning, pages 394–403. PMLR, 2018.

[6] Amos Beimel, Hai Brenner, Shiva Prasad Kasiviswanathan, and Kobbi Nissim. Bounds on the sample complexity for private learning and private data release. Machine learning, 94:401–437, 2014.

[7] Tony Cai, Jianqing Fan, and Tiefeng Jiang. Distributions of angles in random packing on spheres. Journal of Machine Learning Research, 14(21):1837–1864, 2013.

[8] Gabriel Camano-Garcia. Statistics on Stiefel manifolds. Iowa State University, 2006.

[9] Clément L Canonne, Gautam Kamath, and Thomas Steinke. The discrete gaussian for differential privacy. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, Advances in Neural Information Processing Systems, volume 33, pages 15676–15688. Curran Associates, Inc., 2020.

[10] George Casella and Roger L Berger. Statistical inference, volume 2. Duxbury Pacific Grove, CA, 2002.

[11] Thee Chanyaswad, Alex Dytso, H. Vincent Poor, and Prateek Mittal. Mvg mechanism: Differential privacy under matrix-valued query. In Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS '18, pages 230–246, New York, NY, USA, 2018. ACM.

[12] Kamalika Chaudhuri, Anand Sarwate, and Kaushik Sinha. Near-optimal differentially private principal components. Advances in Neural Information Processing Systems, 25:989–997, 2012.

[13] Yasuko Chikuse. Statistics on special manifolds, volume 174. Springer Science & Business Media, 2012.

[14] Harald Cramér. Mathematical methods of statistics, volume 43. Princeton university press, 1999.

[15] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In Serge Vaudenay, editor, Advances in Cryptology - EUROCRYPT 2006, pages 486–503, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

[16] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3, pages 265–284. Springer, 2006.

[17] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. Foundations and Trends® in Theoretical Computer Science, 9(3–4):211–407, 2014.

[18] Cynthia Dwork and Guy N Rothblum. Concentrated differential privacy. arXiv preprint arXiv:1603.01887, 2016.

[19] Cynthia Dwork, Guy N Rothblum, and Salil Vadhan. Boosting and differential privacy. In 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, pages 51–60. IEEE, 2010.

[20] David H Fremlin. Measure theory, volume 4. Torres Fremlin, 2000.

[21] Arjun K Gupta and Daya K Nagar. Matrix variate distributions. Chapman and Hall/CRC, 2018.

[22] Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In Proceedings of the forty-second ACM symposium on Theory of computing, pages 705–714, 2010.

[23] Kaifeng Jiang, Dongxu Shao, Stéphane Bressan, Thomas Kister, and Kian-Lee Tan. Publishing trajectories with differential privacy guarantees. In Proceedings of the 25th International Conference on Scientific and Statistical Database Management, pages 1–12, 2013.

[24] Wuxuan Jiang, Cong Xie, and Zhihua Zhang. Wishart mechanism for differentially private principal components analysis. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 30, 2016.

[25] Norman L Johnson, Samuel Kotz, and Narayanaswamy Balakrishnan. Continuous univariate distributions, volume 2, volume 289. John wiley & sons, 1995.

[26] Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? SIAM Journal on Computing, 40(3):793–826, 2011.

[27] Chinubhai G Khatri. Some results for the singular normal multivariate regression models. Sankhyā: The Indian Journal of Statistics, Series A, pages 267–280, 1968.

[28] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. The cifar-10 dataset. online: http://www. cs. toronto. edu/kriz/cifar. html, 55(5), 2014.

[29] Koon-Shing Kwong and Boris Iglewicz. On singular multivariate normal distribution and its applications. Computational statistics and data analysis, 22(3):271–285, 1996.

[30] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11):2278–2324, 1998.

[31] Chao Li, Michael Hay, Gerome Miklau, and Yue Wang. A data-and workload-aware algorithm for range queries under differential privacy. Proceedings of the VLDB Endowment, 7(5), 2014.

[32] Changchang Liu, Supriyo Chakraborty, and Prateek Mittal. Dependence makes you vulnberable: Differential privacy under dependent tuples. In NDSS, volume 16, pages 21–24, 2016.

[33] Fang Liu. Generalized gaussian mechanism for differential privacy. IEEE Transactions on Knowledge and Data Engineering, 31(4):747–756, 2018.

[34] Aleksandar Nikolov, Kunal Talwar, and Li Zhang. The geometry of differential privacy: the sparse and approximate cases. In Proceedings of the forty-fifth annual ACM symposium on Theory of computing, pages 351–360, 2013.

[35] Lu Ou, Zheng Qin, Shaolin Liao, Yuan Hong, and Xiaohua Jia. Releasing correlated trajectories: Towards high utility and optimal differential privacy. IEEE Transactions on Dependable and Secure Computing, 17(5):1109–1123, 2018.

[36] TensorFlow Privacy. TensorFlow Privacy: MNIST DP-SGD Tutorial. https://github.com/tensorflow/privacy/blob/master/tutorials/mnist_dpsgd_tutorial.py.

[37] Serge B Provost and AM Mathai. Quadratic forms in random variables: theory and applications. M. Dekker, 1992.

[38] Wahbeh Qardaji, Weining Yang, and Ninghui Li. Understanding hierarchical methods for differentially private histograms. Proceedings of the VLDB Endowment, 6(14):1954–1965, 2013.

[39] Calyampudi Radhakrishna Rao and Mathematischer Statistiker. Linear statistical inference and its applications, volume 2. Wiley New York, 1973.

[40] Horn Roger and R Johnson Charles. Topics in matrix analysis, 1994.

[41] Kathrin Schacke. On the kronecker product. Master's thesis, University of Waterloo, 2004.

[42] Muni Shanker Srivastava and CG Khatri. An introduction to multivariate statistics. 2009.

[43] Meng Sun and Wee Peng Tay. On the relationship between inference and data privacy in decentralized iot networks. IEEE Transactions on Information Forensics and Security, 15:852–866, 2019.

[44] TensorFlow. Convolutional Neural Network (CNN). https://www.tensorflow.org/tutorials/images/cnn.

[45] Peter H Westfall. Kurtosis as peakedness, 1905–2014. rip. The American Statistician, 68(3):191–195, 2014.

[46] Yonghui Xiao and Li Xiong. Protecting locations with differential privacy under temporal correlations. In Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, pages 1298–1309, 2015.

[47] Jungang Yang, Liyao Xiang, Jiahao Yu, Xinbing Wang, Bin Guo, Zhetao Li, and Baochun Li. Matrix gaussian mechanisms for differentially-private learning. IEEE Transactions on Mobile Computing, 2021.

[48] Jun Zhao, Teng Wang, Tao Bai, Kwok-Yan Lam, Zhiying Xu, Shuyu Shi, Xuebin Ren, Xinyu Yang, Yang Liu, and Han Yu. Reviewing and improving the gaussian mechanism for differential privacy. arXiv preprint arXiv:1911.12060, 2019.

A Proof of (e) in Equation (5)

Define $\sigma(\mathbf{\Sigma}) \overset{\text{def}}{=} [a_1, a_2, \ldots, a_M]$ and $\sigma(\mathbf{\Psi}) \overset{\text{def}}{=} [b_1, b_2, \ldots, b_N]$. Since $\mathbf{\Sigma}, \mathbf{\Psi} \in \mathbb{PD}$ and they are both symmetric, we have $a_m > 0, \forall m \in [1, M]$ and $b_n > 0, \forall n \in [1, N]$. Then, to prove (e) in (5), we essentially need to show that

$$\sqrt{\sum_{m=1}^{M} a_m^2} \times \sqrt{\sum_{n=1}^{N} b_n^2} \geq \frac{M}{\sqrt{\sum_{m=1}^{M} \frac{1}{a_m^2}}} \times \frac{N}{\sqrt{\sum_{n=1}^{M} \frac{1}{b_n^2}}}. \tag{17}$$

Apply harmonic mean-geometric mean inequality to $a_1^2, a_2^2, \ldots, a_M^2$, we have

$$\frac{a_1^2 + a_2^2 + \ldots + a_M^2}{M} \geq \left(a_1^2 \cdot a_2^2 \cdot \ldots \cdot a_M^2\right)^{\frac{1}{M}}. \tag{18}$$

Similarly, for $\frac{1}{a_1^2}, \frac{1}{a_2^2}, \ldots, \frac{1}{a_M^2}$, we have

$$\frac{\frac{1}{a_1^2} + \frac{1}{a_2^2} + \ldots + \frac{1}{a_M^2}}{M} \geq \left(\frac{1}{a_1^2} \cdot \frac{1}{a_2^2} \cdot \ldots \cdot \frac{1}{a_M^2}\right)^{\frac{1}{M}}. \tag{19}$$

Multiplying (18) and (19) gives

$$\frac{a_1^2 + a_2^2 + \ldots + a_M^2}{M} \times \frac{\frac{1}{a_1^2} + \frac{1}{a_2^2} + \ldots + \frac{1}{a_M^2}}{M}$$
$$\geq \left(a_1^2 \cdot a_2^2 \cdot \ldots \cdot a_M^2\right)^{\frac{1}{M}} \left(\frac{1}{a_1^2} \cdot \frac{1}{a_2^2} \cdot \ldots \cdot \frac{1}{a_M^2}\right)^{\frac{1}{M}} = 1.$$

Taking square root of the above, we have

$$\frac{\sqrt{\sum_{m=1}^{M} a_m^2}}{\sqrt{M}} \times \frac{\sqrt{\sum_{m=1}^{M} \frac{1}{a_m^2}}}{\sqrt{M}} \geq 1 \Leftrightarrow \sqrt{\sum_{m=1}^{M} a_m^2} \geq \frac{M}{\sqrt{\sum_{m=1}^{M} \frac{1}{a_m^2}}}.$$

By applying the same procedure on $[b_1, b_2, \ldots, b_N]$, we have $\sqrt{\sum_{n=1}^{N} b_n^2} \geq \frac{N}{\sqrt{\sum_{n=1}^{N} \frac{1}{b_n^2}}}$. Thus, (17) is proved, and the step (e) in (5) follows.

B $\mathbb{E}[\mathcal{L}]$ versus M for Various Mechanisms

Next, we empirically investigate the accuracy loss of the query results $f(\boldsymbol{x}) \in \mathbb{R}^M$ when M increases. In Figure 9, by assuming $\Delta_2 f = 1$, varying M from 10^2 to 10^6, and fixing $\delta = 10^{-10}$, we plot the expected accuracy loss introduced by the R1SMG, classic Gaussian, and analytic Gaussian mechanisms when $\varepsilon \in \{10^{-7}, 10^{-8}, 10^{-9}\}$. Note that we do not include MVG mechanism in the comparison, because (i) it requires the prior knowledge of $\gamma = \sup_{\boldsymbol{x}} ||f(\boldsymbol{x})||_F$, which depends on specific datasets and applications (e.g., see Theorem 3), and (ii) it will introduce more noise than the analytic Gaussian mechanism.

From Figure 9, we observe that, for both classic and analytic Gaussian mechanisms, the expected accuracy loss (i.e., $Tr[\sigma^2 \mathbf{I}_{M\times M}]$ and $Tr[\sigma_A^2 \mathbf{I}_{M\times M}]$) scales linearly with M, which validates our theoretical findings in Section 3. In contrast,

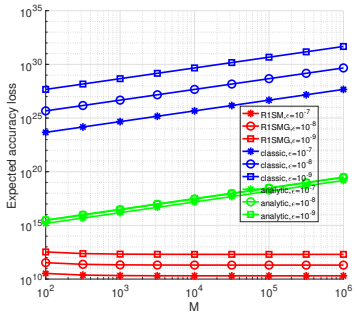

Figure 9: Expected accuracy loss of the R1SMG, classic Gaussian, and analytic Gaussian mechanisms under increasing query size when $\varepsilon \in \{10^{-7}, 10^{-8}, 10^{-9}\}$ and $\delta = 10^{-10}$.

as M increases, the expected accuracy loss incurred by the R1SMG mechanism (i.e., $Tr[\mathbf{\Pi}] = \sigma_*$) decreases and converges to $\frac{2(\Delta_2 f)^2}{\varepsilon}$. Thus, the R1SMG mechanism not only breaks the curse of full-rank covariance matrices, but also significantly improves the accuracy (utility) of the perturbed $f(\boldsymbol{x})$, as it is able to achieve (ε, δ)-DP by using much weaker noise even when the dimension of $f(\boldsymbol{x})$ increases.

C Potential Improvement of the MVG

When studying the MVG mechanism, we observe one potential solution to improve its privacy guarantee. In particular, one intermediate step in the proof of the MVG mechanism requires deriving an upper bound for

$$\clubsuit = Tr\Big[\mathbf{\Psi}^{-1}\mathbf{Y}^T\mathbf{\Sigma}^{-1}\Delta + \mathbf{\Psi}^{-1}\Delta^T\mathbf{\Sigma}^{-1}\mathbf{Y} + \mathbf{\Psi}^{-1}f(D_2)^T\mathbf{\Sigma}^{-1}f(D_2) - \mathbf{\Psi}^{-1}f(D_1)^T\mathbf{\Sigma}^{-1}f(D_1)\Big], \tag{20}$$

where $\mathbf{Y}$ stands for the output of the MVG mechanism and $\Delta = f(D_1) - f(D_2)$ (see [11] page 244, right column). Due to the negative sign in the fourth term, the authors simply bound its absolute value, i.e., $\left|Tr\left[\mathbf{\Psi}^{-1}f(D_1)^T\mathbf{\Sigma}^{-1}f(D_1)\right]\right|$, which results in a very loose upper bound related to the quadratic form in the Frobenius norm of the matrix query result (i.e, γ^2 in Theorem 3), and further compromises the utility.

In fact, the original proof of the MVG mechanism can be improved if we apply the transformation

$$Tr\left[\mathbf{\Psi}^{-1}f(D_2)^T\mathbf{\Sigma}^{-1}f(D_2) - \mathbf{\Psi}^{-1}f(D_1)^T\mathbf{\Sigma}^{-1}f(D_1)\right] =$$
$$Tr\left[\mathbf{\Psi}^{-1}\big(f(D_2)^T\mathbf{\Sigma}^{-1}\big(f(D_2) - f(D_1)\big) + \big(f(D_2) - f(D_1)\big)^T\mathbf{\Sigma}^{-1}f(D_1)\big)\right],$$

and rewrite $\clubsuit$ as

$$\clubsuit = Tr\Big[\mathbf{\Psi}^{-1}\big(\mathbf{Y} - f(D_2)\big)^T\mathbf{\Sigma}^{-1}\Delta + \mathbf{\Psi}^{-1}\Delta^T\mathbf{\Sigma}^{-1}\big(\mathbf{Y} - f(D_1)\big)\Big].$$

Then, we can bound $\clubsuit$ using singular values, and hence the quadratic form in the Frobenius norm of the query result can be completely removed. We do not follow the above steps to further develop an improved version of the MVG mechanism, because the improved version is still a "victim" of the identified curse of full-rank covariance matrices.