Calibrating Noise for Group Privacy in Subsampled Mechanisms

Yangfan Jiang
National University of Singapore
jyangfan@comp.nus.edu.sg

Xinjian Luo
National University of Singapore
xinjluo@comp.nus.edu.sg

Yin Yang
College of Sci. and Engr., Hamad Bin Khalifa University
yyang@hbku.edu.qa

Xiaokui Xiao
National University of Singapore
xkxiao@nus.edu.sg

ABSTRACT

Given a group size m and a sensitive dataset D, group privacy (GP) releases information about D (e.g., weights of a neural network trained on D) with the guarantee that the adversary cannot infer with high confidence whether the underlying data is D or a neighboring dataset D' that differs from D by m records. GP generalizes the well-established notion of differential privacy (DP) for protecting individuals' privacy; in particular, when $m = 1$, GP reduces to DP. Compared to DP, GP is capable of protecting the sensitive *aggregate* information of a group of up to m individuals, e.g., the average annual income among members of a yacht club. Despite its longstanding presence in the research literature and its promising applications, GP is often treated as an afterthought, with most approaches first developing a differential privacy (DP) mechanism and then using a generic conversion to adapt it for GP, treating the DP solution as a black box. As we point out in the paper, this methodology is suboptimal when the underlying DP solution involves subsampling, e.g., in the classic DP-SGD method for training deep learning models. In this case, the DP-to-GP conversion is overly pessimistic in its analysis, leading to high error and low utility in the published results under GP.

Motivated by this, we propose a novel analysis framework that provides tight privacy accounting for subsampled GP mechanisms. Instead of converting a black-box DP mechanism to GP, our solution carefully analyzes and utilizes the inherent randomness in subsampled mechanisms, leading to a substantially improved bound on the privacy loss with respect to GP. The proposed solution applies to a wide variety of foundational mechanisms with subsampling. Extensive experiments with real datasets demonstrate that compared to the baseline convert-from-blackbox-DP approach, our GP mechanisms achieve noise reductions of over an order of magnitude in several practical settings, including deep neural network training.

PVLDB Reference Format:
Yangfan Jiang, Xinjian Luo, Yin Yang, and Xiaokui Xiao. Calibrating Noise for Group Privacy in Subsampled Mechanisms. PVLDB, 18(2): 322 - 334, 2024.
doi:10.14778/3705829.3705848

PVLDB Artifact Availability:
The source code, data, and/or other artifacts have been made available at https://github.com/Yangfan-Jiang/calibrating-group-privacy.

Proceedings of the VLDB Endowment, Vol. 18, No. 2 ISSN 2150-8097.
doi:10.14778/3705829.3705848

1 INTRODUCTION

With the rapid advances of machine learning techniques, data privacy has become a growing concern, and simple measures often fail to provide adequate protection to prevent leakage of sensitive information [25, 45]. Differential privacy (DP) [19, 20] is a strong and rigorous standard for ensuring individuals' privacy, which has gained adoption in industry [4, 14, 22] and widespread interest in academia [1, 7, 8, 15, 17, 32, 36, 39, 60]. In many practical scenarios, however, safeguarding only individual-level information may be insufficient, as aggregates over a group of individuals can also be highly sensitive [26, 27, 30, 31, 38, 47]. For instance, the income distribution of a bank's private banking customers can be a critical business secret. To tackle this issue, one natural approach is to extend the notion of DP to *group privacy* (*GP*), which protects the aggregate information of a group of individuals.

Specifically, a randomized algorithm $\mathcal{A}$ satisfies GP with group size m if, for any pair of neighbor datasets D and D' differing by m records, the output distributions of $\mathcal{A}(D)$ and $\mathcal{A}(D')$ are guaranteed to be indistinguishable in an information-theoretic sense, measured by specific privacy parameters, elaborated later in Section 2. In the special case that $m = 1$, this reduces to the classic DP definition. Note that the privacy guarantee in GP indicates that no group of m records can have a significant impact on $\mathcal{A}$'s output distribution; hence, the larger the group size m, the stronger the guarantee. A foundational approach for achieving GP (which includes its special case DP) is to perturb the exact (i.e., non-private) result by injecting a calibrated amount of random noise [16]. Intuitively, a larger group size m, which corresponds to a stronger privacy guarantee, requires a higher amount of injected random noise to satisfy GP, leading to lower result utility.

Since DP is a hot research topic, many well-optimized solutions are available for enforcing DP in various problem settings. To satisfy the more general GP requirement, a common approach (e.g., [42]) is to convert an existing DP mechanism through a generic conversion procedure. Intuitively, since DP is a special case of GP with $m = 1$, we can obtain a GP-compliant mechanism by scaling the noise injected by DP by a factor that depends on the value of m. Note that here, the conversion algorithm is generic, meaning that it treats the underlying DP mechanism as a black box without considering the unique properties of the problem or the DP solution.

In this paper, we focus on the problem of releasing analysis results under GP, where the analysis involves random *subsampling*,

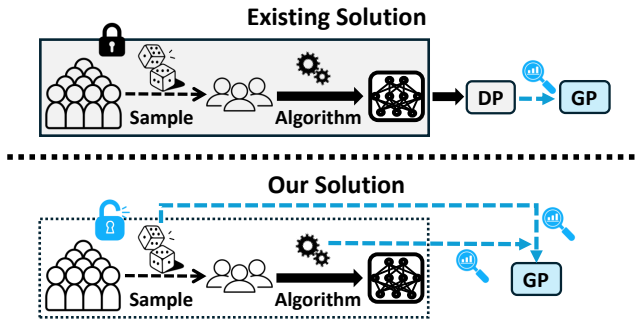

Figure 1: High-level idea: the existing solution converts a black-box DP mechanism, while our approach conducts a direct, white-box-style privacy analysis.

as follows. Given an input dataset D and an analysis algorithm $\mathcal{F}$, each record in D is randomly selected into a subset S with a probability of q ($0 < q < 1$); after that, the analysis result is obtained by performing $\mathcal{F}$ on the subset S instead of the entire dataset D. A notable example of subsampled analysis is stochastic gradient descent (SGD), which is commonly used to train large-scale machine learning models such as neural networks. It has been shown that for such applications, the amount of noise required to satisfy DP can be significantly reduced through sophisticated privacy accounting methods that exploit properties of subsampling [1, 5, 43, 51, 61], e.g., in the classic DP-SGD algorithm [1] for deep learning with DP. The intuition is that subsampling already provides a certain level of privacy protection, in the sense that as long as an individual record is not included in the sample set S, no information about this record is leaked even in the exact result $\mathcal{F}(S)$. This inherent privacy protection is amplified as the sampling rate q becomes lower.

Main observation. To enforce GP with a given group size m on subsampled analysis results, a baseline approach would be to take the corresponding DP mechanism and apply the above-mentioned generic conversion, by scaling the random noise by a factor determined by m. The problem with this approach is that the generic conversion, clarified in Section 2.2, treats the underlying DP mechanism as a black box and, thus, fails to utilize the inherent randomness introduced by subsampling, leading to suboptimal result utility. To illustrate, consider two neighbor datasets D and D' that differ by m records, and a subsampling analysis that selects each record with probability q. Let S and S' be the sample sets obtained from D and D', respectively; for simplicity, let's consider the case that the same random seed is used to obtain both S and S'. Then, intuitively, S and S' should differ by qm records in expectation, rather than m, which has a vanishingly low probability of $q^m \ll 1$. To be more precise, the number of different records between S and S' follows the binomial distribution $B(m, q)$, with mean qm and variance $q(1-q)m$. When q is sufficiently low, e.g., $q = O(1/m)$, the variance of $B(m, q)$ becomes a constant, in which case the number of different records between S and S' is tightly concentrated around its mean value qm. This hints that we might be able to achieve GP by scaling the noise injected by DP by a factor that depends on qm, which is significantly smaller than m as in the baseline approach.

Our contributions. In this paper, we establish a refined group privacy bound for subsampled mechanisms through a more sophisticated and precise privacy analysis. Our analysis follows the framework of Rényi group privacy (RGP) [42], which enables more accurate privacy guarantees for subsampled and iterative mechanisms such as DP-SGD [1]. Further, an RGP mechanism can be transformed to satisfy traditional notions of GP, as elaborated in Section 2.1. Unlike existing methods that simply convert a black-box DP mechanism, our approach is subsampling-aware, which directly analyzes the Rényi divergence between the output distributions of the subsampled analysis on pairs of datasets that differ by m records, as illustrated in Figure 1. This direct, white-box-style analysis offers several advantages for tightening up the RGP guarantee: (i) it accounts for specific algorithmic characteristics, such as the shape of the output distribution; (ii) it harnesses the inherent randomness in subsampled mechanisms for amplifying the privacy guarantee; and (iii) it significantly reduces the impact of worst-case scenarios on group privacy guarantees. These properties help significantly enhance our RGP guarantee for subsampled mechanisms.

Through rigorous theoretical analysis, we prove that our bound, in general, offers a substantially improved RGP guarantee compared to previous methods for the subsampled mechanisms. For instance, in the case of the subsampled Gaussian mechanism [43], a core component in many widely-used privacy-preserving algorithms, including DP-SGD and its adaptations [2, 7, 26, 56, 61], our bound leads to noise reduction by a multiplicative factor of approximately $O(m^{0.58})$ compared to existing methods, where m is the group size. Further, we prove the tightness of our general RGP bound for subsampled mechanisms. Specifically, we first establish an analytical lower bound of RGP guarantee for subsampled mechanisms by constructing a carefully crafted pair of neighboring datasets. Then, we show that our proposed RGP analysis asymptotically matches this lower bound, thereby justifying the tightness of our RGP bound.

Besides asymptotic improvements, the proposed RGP analysis framework has practical implications across various applications, in the sense that concrete instantiations of our RGP bound can be employed to derive significantly improved RGP guarantees for various privacy-preserving algorithms. Specifically, we present exact, *closed-form* RGP bounds for several widely-used mechanisms, including the subsampled Gaussian, Laplace, Skellam (which is often used to enforce DP in secure federated learning [2, 7]), and Randomized Response (commonly used in local DP [22]) mechanisms.

A result that might be of independent interest is that our analysis for the subsampled Laplace mechanism is not only a significant improvement for the RGP guarantee but for the popular Rényi differential privacy (RDP) definition as well. Specifically, for the d-dimensional Laplace mechanism, existing RDP methods, to our knowledge, all involve privacy composition, which accumulates privacy costs in each dimension. In other words, this approach incurs a multiplicative factor of $\Omega(d)$ in the privacy cost, which is prohibitively high when d is large. To tackle this issue, we formulate the task of deriving a privacy guarantee for the d-dimensional Laplace mechanism as a constrained optimization problem, and demonstrate that this problem can be simplified to a more manageable one-dimensional scenario, thereby avoiding the composition.

Finally, we conduct a thorough comparison of the proposed GP bound with existing methods through numerical experiments. The

results demonstrate a significant improvement of our method over existing ones, typically by over an order of magnitude in terms of injected error scale. Our results also validate the tightness of our RGP bound, which closely matches the theoretical lower bound under a wide range of configurations. In addition, we apply our results to enforcing GP on SGD for deep neural network training, using the MNIST, Fashion-MNIST, and CIFAR-10 benchmark datasets. The results show that compared to existing methods, our solution consistently achieves considerably higher model utility across different group sizes and privacy parameter settings.

2 PRELIMINARIES

2.1 Group Privacy

First, we formally define the notion of neighbor datasets, which is a major building block of the GP guarantee, as follows.

Definition 2.1 (m-neighboring datasets)**.** *Two datasets D and D' are m-neighboring datasets if and only if they differ by m records.*

In the above definition, parameter m is referred to as the *group size*. When m is clear from the context, we simply call D and D' neighbor databases. Note that there are two cases that D and D' differ by m-records: (i) D' can be obtained by adding or removing m records from D, which is referred to as the *unbounded* definition [18, 29], and (ii) D' can be obtained by replacing m existing records in D, which is called the *bounded* definition [20, 29]. Accordingly, GP defined using the bounded (resp., unbounded) definition of neighbor databases is referred to as bounded (resp, unbounded) GP.

Next we present the classic definition of GP, as follows.

Definition 2.2 ((m, ϵ, δ)-Group Privacy [21, 49])**.** *A randomized algorithm $\mathcal{A}$ satisfies (m, ϵ, δ)-DP if, for any two m-neighboring datasets D, D', and for any subset of possible outputs $O \subseteq Range(\mathcal{A})$, it holds that*

$$\Pr\left[\mathcal{A}\left(D\right) \in O\right] \leq e^{\epsilon} \Pr\left[\mathcal{A}\left(D'\right) \in O\right] + \delta. \tag{1}$$

A notable special case is that when $m = 1$, (m, ϵ, δ)-GP reduces to (ϵ, δ)-differential privacy [19, 20]. The parameters ϵ and δ control the trade-off between privacy and utility. Smaller values of ϵ and δ result in more similar output distributions of $\mathcal{A}$ on the neighbor datasets D and D', thereby providing stronger privacy protection. Meanwhile, as mentioned in Section 1, the group size m also affects the GP guarantee: a larger m prevents leakage of information (both individual and aggregate) derived from a larger group of individuals, which corresponds to a stronger guarantee, and vice versa.

Note that by definition, only a randomized algorithm can satisfy GP. Given a deterministic algorithm $\mathcal{F}$, we can ensure GP by injecting random noise to $\mathcal{F}$'s output [16], where the noise magnitude scales inversely with the privacy parameters ϵ and δ, and at the same time positively correlated with the group size m [20].

Rényi group privacy (RGP) [42]. RGP is an alternative notion of group privacy that measures the indistinguishability between outputs of a randomized algorithm $\mathcal{A}$ using Rényi divergence [44]. RGP is particularly effective in accurately tracking cumulative privacy costs in iterative and subsampled mechanisms, due to its strong privacy composition property [42, 43], explained shortly. For this reason, the Rényi differential privacy (which is a special case of RGP with group size $m = 1$) plays a pivotal role in the analyses of the DP-SGD algorithm [1, 43] as well as its variants [2, 7, 26, 56, 61]. Formally, Rényi divergence and RGP are defined as follows.

Definition 2.3 (Rényi Divergence [44, 50])**.** *Given two probability distributions P and Q that are defined on the same probability space $\mathcal{Z}$, let $P(z)$ and $Q(z)$ denote the densities of P and Q at $z \in \mathcal{Z}$. The Rényi divergence of a finite order $\alpha > 1$ is*

$$D_\alpha(P\|Q) = \frac{1}{\alpha - 1} \log \mathbb{E}_{z\sim Q}\left[\left(\frac{P(z)}{Q(z)}\right)^{\alpha}\right]. \tag{2}$$

Definition 2.4 ((m, α, τ)-Rényi Group Privacy [42])**.** *A randomized algorithm $\mathcal{A}$ is said to satisfy (m, α, τ)-Rényi group privacy (RGP), if for any pair of m-neighboring datasets D and D', we have*

$$D_\alpha(\mathcal{A}(D)\|\mathcal{A}(D')) \leq \tau,$$

where $D_\alpha(\cdot\|\cdot)$ is the Rényi divergence of order α.

When the group size $m = 1$, the (m, α, τ)-RGP is referred to as (α, τ)-Rényi differential privacy (RDP). The following lemma presents the sequential composition property of RGP. As mentioned above, this property is particularly useful for iterative mechanism design (e.g., in DP-SGD [1] and its variants), in which privacy cost accumulates across the iterations.

Lemma 2.1 (Sequential Composition of RGP [42])**.** *Let $\mathcal{M}_1 : \mathcal{D} \mapsto \mathcal{R}_1$ and $\mathcal{M}_2 : \mathcal{R}_1 \times \mathcal{D} \mapsto \mathcal{R}_2$ be two randomized mechanisms with independent source of randomness that satisfy (m, α, τ_1) and (m, α, τ_2)-RGP, respectively. Then the combination of these two mechanisms, defined by $\mathcal{M}_{1,2}(D) := (\mathcal{M}_1(D), \mathcal{M}_2(\mathcal{M}_1(D), D))$, satisfies $(m, \alpha, \tau_1 + \tau_2)$-RGP.*

Finally, an (m, α, τ)-RGP can be converted into an (m, ϵ, δ)-GP guarantee for any $\delta \in (0, 1)$ through the following lemma:

Lemma 2.2 (From RGP to GP [6, 12])**.** *If a randomized algorithm $\mathcal{A}$ satisfies (m, α, τ)-RGP, then for all $\delta \in (0, 1)$, the algorithm $\mathcal{A}$ also satisfies (m, ϵ, δ)-GP, where*

$$\epsilon = \tau + \frac{\log(1/\delta) + (\alpha - 1)\log(1 - 1/\alpha) - \log\alpha}{\alpha - 1}.$$

In practice, given a δ and an m, we determine the final (m, ϵ, δ)-group privacy guarantee by minimizing ϵ over α using Lemma 2.2. Typically, α ranges from 2 to 100, as commonly seen in production-ready libraries such as TensorFlow Privacy[1], Opacus[2], and autodp[3].

2.2 From RDP to RGP

We now introduce the existing general methodology for converting RDP guarantees to RGP guarantees.

Lemma 2.3 (From RDP to RGP [42])**.** *Let $c \in \mathbb{N}$ be an arbitrary positive integer. If $\mathcal{A}$ satisfies (α, τ)-RDP and $\alpha \geq 2^{c+1}$, then*

$$D_{\alpha/2^c}(D\|D') \leq 3^c\tau,$$

for all pairs of 2^c-neighboring datasets D and D'.

Accordingly, to convert a (α, τ)-RDP to RGP with a given group size m, we can first find $c \in \mathbb{N}$ such that $2^c \geq m$, and then scale down α to obtain $\alpha' = \frac{\alpha}{2^c}$, and at the same time scale up τ to

[1] https://github.com/tensorflow/privacy
[2] https://github.com/pytorch/opacus
[3] https://github.com/yuxiangw/autodp

$\tau' = 3^c\tau$, to obtain the corresponding (m, α', τ')-RGP guarantee. Observe that since the above lemma is only applicable for $c \in \mathbb{N}$, the conversion needs to increase m to the next power of two. Since τ is scaled up by a factor of 3^c which is approximately $m^{1.58}$ (since $\frac{\log(3)}{\log(2)} \approx 1.58$), this leads to a rather conservative bound on privacy cost, and, thus, high error scale required to satisfy RGP. As we shall elaborate in Sections 3 and 4, our results can derive (m, α, τ)-RGP guarantees for any $m \in \mathbb{N}$, overcoming this issue.

Lastly, we mention that references [21, 49] in which GP is defined also include a basic DP-to-GP conversion method that transforms a (ϵ', δ')-DP guarantee to (m, ϵ, δ)-GP, where $\epsilon = m\epsilon'$ and $\delta = \frac{e^{m\epsilon'}-1}{e^{\epsilon'}-1}\delta'$. Note that as ϵ' grows, δ approaches $e^m\delta'$, which can become excessively large even with a moderate group size m, e.g., 64. In practice, it is often required that δ is no more than $o(1/n)$ where n is the number of records in the underlying dataset, since otherwise, one can release the exact value of a random record while still satisfying GP. Further, as mentioned before, the (m, ϵ, δ)-GP definition does not have the nice composition properties of RGP (i.e., Lemma 2.1) and, thus, is not easy to use for iterative algorithms such as SGD. Hence, this paper focuses on the RGP definition [42] and the corresponding RDP-to-RGP conversion rule described above.

2.3 Subsampled Mechanism

Let D be a dataset and let $\mathcal{A}$ be a randomized algorithm. We adopt the definition of the subsampled mechanism used in the DP-SGD algorithm [1, 43], as follows.

Definition 2.5 (Subsampled Mechanism)**.** *Given an input dataset D, the subsampled mechanism constructs a subset $S \subseteq D$ by including each record $x \in D$ into S independently with a fixed probability of $q \in (0, 1)$. $\mathcal{A}$ is then performed on S to produce the privatized output. Formally, this is defined as:*

$$\mathcal{M}(D) \triangleq (\mathcal{A} \circ \mathsf{Subsample})(D) = \mathcal{A}(\mathsf{Subsample}(D)),$$

where $\mathsf{Subsample}(D)$ *denotes the subsampling procedure that constructs the subset S from D.*

The output distribution of $\mathcal{M}$ is essentially a mixture distribution, where the distributions of $\mathcal{A}(S)$ for all $S \subseteq D$ are the mixed components, and their corresponding mixture weights are the probabilities of S being sampled. In other words, the distribution of $\mathcal{M}(D)$ can be expressed as:

$$\mathcal{M}(D) \sim \sum_{S \subseteq D} p_S \mathcal{A}(S),$$

where p_S denotes the probability of constructing the subset $S \subseteq D$. Note that in the literature, different mechanisms may use different subsampling procedures. For instance, in DPIS [53], each record is selected into the sample set S based on a probability computed based on an importance measure of the record. Another approach [51] involves uniformly sampling a subset from all possible subsets of a predetermined set size. In this paper, we focus on the most widely-used subsampling procedure, which samples each record independently with a constant probability q. This scheme is used in DP-SGD [1] as well as several of its variants [2, 7, 26, 56, 61]. The analysis of group privacy guarantees on other types of subsampled mechanisms will be considered in future work.

3 MAIN RESULTS

This section presents the main results of the paper: a refined, generic privacy cost upper bound for subsampled mechanisms under RGP. This is essentially a subsampling-aware privacy accounting framework, whose concrete instantiations for specific subsampling mechanisms are presented later in Section 4. The key insight that we utilize for deriving the refined RGP bound is that the subsampling procedure amplifies group privacy. More detailed explanations of this insight behind our analysis are provided in Appendix A of the full version [28].

In what follows, we first introduce this generic RGP bound for subsampled mechanisms in Section 3.1 and derive its corresponding proof in Section 3.2. Then, in Section 3.3, we prove the tightness of our RGP bound.

3.1 General RGP Bound

The following theorem presents our general upper bound on the privacy cost under RGP for subsampled mechanisms.

Theorem 3.1 (RGP upper bound of Subsampled Mechanisms)**.** *Let $\mathcal{M} := \mathcal{A} \circ \mathsf{Subsample}$ be a subsampled mechanism with a sampling rate $q \in (0, 1)$. If $\mathcal{A}$ satisfies $(k, \alpha, \tau_k^*(\alpha))$-*bounded *(resp.* unbounded*) RGP for $k \in \{0, 1, \ldots, m\}$, then $\mathcal{M}$ satisfies $(m, \alpha, \tau_m(\alpha))$-*bounded *(resp.* unbounded*) RGP, where*

$$\tau_m(\alpha) = \frac{1}{\alpha - 1} \log\left(\sum_{k=0}^{m} \binom{m}{k} (1-q)^{m-k} q^k \exp\left((\alpha - 1)\tau_k^*(\alpha)\right)\right).$$

The above theorem generally applies to all subsampled mechanisms with a constant sampling rate. To apply this upper bound to determine the exact RGP guarantee of a specific subsampled mechanism $\mathcal{M} = \mathcal{A} \circ \mathsf{Subsample}$, a practitioner needs to determine the RGP guarantee $\tau_k^*(\alpha)$ of $\mathcal{A}$ for each group size $k \in \{0, 1, \ldots, m\}$. In practice, the algorithm $\mathcal{A}$ is often derived from well-studied DP mechanisms, such as the Gaussian mechanism. Hence, we can usually bypass the existing RDP-to-RGP conversion method (Lemma 2.3), which is rather conservative as explained in Section 2.2, and directly derive a tight group privacy guarantee using properties of these well-studied mechanisms, elaborated in Section 4.

3.2 Proof of Our General RGP Bound

The proof of Theorem 3.1 is established by deriving upper bounds for both $D_\alpha(\mathcal{M}(D) \| \mathcal{M}(D'))$ and $D_\alpha(\mathcal{M}(D') \| \mathcal{M}(D))$, where D and D' are a pair of m-neighboring datasets. In what follows, we first establish the upper bound of the *unbounded* RGP guarantee for $\mathcal{M}$, and then derive the upper bound of the *bounded* RGP guarantee.

Ensuring unbounded RGP. Assume that $\mathcal{A}$ satisfies $(k, \alpha, \tau_k^*(\alpha))$-unbounded RGP for each $k \in \{0, 1, \ldots, m\}$. Without loss of generality, consider D and D' as a pair of unbounded m-neighboring datasets such that $D \subset D'$, with $|D| = n$ and $|D' \setminus D| = m$. Let $\mathcal{B} := \{B \mid B \subseteq D\}$ denote the power set of D, and $\mathcal{J} := \{J \mid J \subseteq D' \setminus D\}$ represent the power set of $D' \setminus D$. Let p_B denote the probability that B is the outcome of the subsampling process $\mathsf{Subsample}(D)$. Since the Subsample procedure places each record into the subset B independently with a constant probability q, the values of p_B are identical for both $\mathcal{M}(D)$ and $\mathcal{M}(D')$. Similarly, denote by p_J the probability that J is subsampled by $\mathcal{M}(D')$ from $D' \setminus D$. Then, the

output distributions of $\mathcal{M}(D)$ and $\mathcal{M}(D')$ can be expressed as

$$\mathcal{M}(D) \sim \sum_{B\in\mathcal{B}} p_B \mathcal{A}(B) \quad \text{and} \quad \mathcal{M}(D') \sim \sum_{B\in\mathcal{B}} p_B \sum_{J\in\mathcal{J}} p_J \mathcal{A}(B\cup J).$$

To establish the RGP guarantee for $\mathcal{M}$, we seek to upper bound $D_\alpha(\mathcal{M}(D)\|\mathcal{M}(D'))$ and $D_\alpha(\mathcal{M}(D')\|\mathcal{M}(D))$ simultaneously. Beginning with the term $D_\alpha(\mathcal{M}(D')\|\mathcal{M}(D))$, we have

$$\begin{aligned} D_\alpha(\mathcal{M}(D')\|\mathcal{M}(D)) =& D_\alpha\left(\sum_{B\in\mathcal{B}} p_B \sum_{J\in\mathcal{J}} p_J \mathcal{A}(B\cup J)\middle\|\sum_{B\in\mathcal{B}} p_B\mathcal{A}(B)\right) \\ \le& \sup_B D_\alpha\left(\sum_{J\in\mathcal{J}} p_J\mathcal{A}(B\cup J)\middle\|\mathcal{A}(B)\right), \end{aligned} \tag{3}$$

where the inequality follows from the joint quasi-convexity of Rényi divergence (see Corollary B.1 in the full version [28]).

Define $\widetilde{B} \coloneqq \arg\max_{B\in\mathcal{B}} D_\alpha\left(\sum_{J\in\mathcal{J}} p_J\mathcal{A}(B\cup J)\middle\|\mathcal{A}(B)\right)$. Denote by $\boldsymbol{\mu}_J$ and $\boldsymbol{\mu}_0$ the probability density functions (pdfs) of $\mathcal{A}(\widetilde{B}\cup J)$ and $\mathcal{A}(\widetilde{B})$, respectively. Then by (3), we can further simplify the upper bound of $D_\alpha(\mathcal{M}(D')\|\mathcal{M}(D))$ as follows:

$$D_\alpha\big(\mathcal{M}(D')\big\|\mathcal{M}(D)\big) \le D_\alpha\left(\sum_{J\in\mathcal{J}} p_J\boldsymbol{\mu}_J\middle\|\boldsymbol{\mu}_0\right). \tag{4}$$

The upper bound of the term $D_\alpha(\mathcal{M}(D)\|\mathcal{M}(D'))$ is derived similarly. Let $\widetilde{B}' \coloneqq \arg\max_{B\in\mathcal{B}} D_\alpha\left(\mathcal{A}(B)\middle\|\sum_{J\in\mathcal{J}} p_J\mathcal{A}(B\cup J)\right)$ and denote by $\boldsymbol{\mu}'_J$ and $\boldsymbol{\mu}'_0$ the pdfs of $\mathcal{A}(\widetilde{B}'\cup J)$ and $\mathcal{A}(\widetilde{B}')$, respectively. It then holds that

$$D_\alpha\big(\mathcal{M}(D)\big\|\mathcal{M}(D')\big) \le D_\alpha\left(\boldsymbol{\mu}'_0\middle\|\sum_{J\in\mathcal{J}} p_J\boldsymbol{\mu}'_J\right). \tag{5}$$

Let A_α and A'_α denote the right hand side (RHS) of (4) and (5), respectively. The subsampled mechanism $\mathcal{M}$ then satisfies $(m,\alpha,\max\{A_\alpha,A'_\alpha\})$-unbounded RGP. We proceed by establishing an upper bound for $\max\{A_\alpha,A'_\alpha\}$. For the term A_α, we have

$$\begin{aligned} A_\alpha &= \frac{1}{\alpha-1}\log\mathbb{E}_{z\sim\boldsymbol{\mu}_0}\left[\left(\frac{\sum_{J\in\mathcal{J}} p_J\boldsymbol{\mu}_J(z)}{\boldsymbol{\mu}_0(z)}\right)^\alpha\right] \\ &= \frac{1}{\alpha-1}\log\int_{\mathcal{Z}} \frac{\left(\sum_{J\in\mathcal{J}} p_J\boldsymbol{\mu}_J(z)\right)^\alpha}{\boldsymbol{\mu}_0(z)^{\alpha-1}}\mathrm{d}z \\ &\le \frac{1}{\alpha-1}\log\left(\sum_{J\in\mathcal{J}} p_J\int_{\mathcal{Z}}\frac{\boldsymbol{\mu}_J(z)^\alpha}{\boldsymbol{\mu}_0(z)^{\alpha-1}}\mathrm{d}z\right) \\ &= \frac{1}{\alpha-1}\log\left(\sum_{J\in\mathcal{J}} p_J\exp\big((\alpha-1)D_\alpha(\boldsymbol{\mu}_J\|\boldsymbol{\mu}_0)\big)\right), \end{aligned} \tag{6}$$

where the inequality follows from the convexity of the function $f(x)=x^\alpha$ for all $\alpha>1$, and the last equality is derived from the definition of Rényi divergence.

Let $\mathcal{J}_k \coloneqq \{J\in\mathcal{J} \mid |J|=k\}$. Because $\mathcal{A}$ satisfies $(k,\alpha,\tau_k^*(\alpha))$-unbounded RGP, it is established that $D_\alpha\big(\boldsymbol{\mu}_J\big\|\boldsymbol{\mu}_0\big) \le \tau_k^*(\alpha)$ for all $J\in\mathcal{J}_k$. Consequently, it follows that

$$\begin{aligned} (6) =& \frac{1}{\alpha-1}\log\left(\sum_{k=0}^m\sum_{J\in\mathcal{J}_k} p_J\exp\big((\alpha-1)D_\alpha(\boldsymbol{\mu}_J\|\boldsymbol{\mu}_0)\big)\right) \\ \le& \frac{1}{\alpha-1}\log\left(\sum_{k=0}^m p_k\exp\Big((\alpha-1)\tau_k^*(\alpha)\Big)\right), \end{aligned} \tag{7}$$

where $p_k \coloneqq \sum_{J\in\mathcal{J}_k} p_J$.

We now proceed to upper bound the term A'_α. Because Rényi divergence is convex in its second term (see Lemma B.2 in [28]), we immediately obtain

$$A'_\alpha = D_\alpha\left(\boldsymbol{\mu}'_0\middle\|\sum_J p_J\boldsymbol{\mu}'_J\right) \le \sum_J p_J D_\alpha\left(\boldsymbol{\mu}'_0\middle\|\boldsymbol{\mu}'_J\right) \le \sum_{k=0}^m p_k\tau_k^*(\alpha). \tag{8}$$

Denote by $\overline{A_\alpha}$ and $\overline{A'_\alpha}$ the RHS of (7) and (8), respectively. Given the concavity of the logarithm function, we have

$$\begin{aligned} \overline{A_\alpha} &= \frac{1}{\alpha-1}\log\left(\sum_{k=0}^m p_k\exp\Big((\alpha-1)\tau_k^*(\alpha)\Big)\right) \\ &\ge \frac{1}{\alpha-1}\sum_{k=0}^m p_k\log\Big(\exp\Big((\alpha-1)\tau_k^*(\alpha)\Big)\Big) = \sum_{k=0}^m p_k\tau_k^*(\alpha) = \overline{A'_\alpha}. \end{aligned}$$

Therefore, $\max\{A_\alpha,A'_\alpha\}$ can be upper bounded as:

$$\begin{aligned} \max\{A_\alpha,A'_\alpha\} &\le \max\{\overline{A_\alpha},\overline{A'_\alpha}\} \\ &= \overline{A_\alpha} = \frac{1}{\alpha-1}\log\left(\sum_{k=0}^m p_k\exp\Big((\alpha-1)\tau_k^*(\alpha)\Big)\right). \end{aligned} \tag{9}$$

Recall that in the underlying subsampled mechanism, each record in $D'\setminus D$ is randomly selected into the input subset independently with a fixed probability q. Accordingly, the number of records in J follows a binomial distribution with m trials and success probability q. Hence, we have

$$p_k = \binom{m}{k}(1-q)^{m-k}q^k. \tag{10}$$

Substituting (10) into (9) yields the privacy guarantee with respect to unbounded RGP.

Ensuring *bounded* RGP. We slightly abuse the notations and let D and D' be a pair of m-bounded neighboring datasets with $|D|=|D'|=n$. Without loss of generality, we assume that D and D' differ in the last m records. Suppose $\mathcal{A}$ satisfies $(k,\alpha,\tau_k^*(\alpha))$-bounded RGP for each $k\in\{0,1,\dots,m\}$. Define $\mathcal{I} \coloneqq \{I \mid I\subseteq\{1,2,\dots,n\}\}$ as the power set of the indices $\{1,2,\dots,n\}$ and let $D_I \coloneqq \{\boldsymbol{x}_i \mid i\in I, \boldsymbol{x}_i\in D\}$ be the set of records in D indexed by I. Denote by p_I the probability that the subset D_I is the outcome of Subsample(D). Consequently, the output distributions of $\mathcal{M}(D)$ and $\mathcal{M}(D')$ can be expressed as:

$$\mathcal{M}(D) \sim \sum_{I\in\mathcal{I}} p_I\mathcal{A}(D_I) \quad \text{and} \quad \mathcal{M}(D') \sim \sum_{I\in\mathcal{I}} p_I\mathcal{A}(D'_I).$$

Let ν_I and ν'_I denote the pdfs of $\mathcal{A}(D_I)$ and $\mathcal{A}(D'_I)$. Then

$$\begin{aligned} D_\alpha(\mathcal{M}(D)\|\mathcal{M}(D')) =& \frac{1}{\alpha-1}\log\int_{\mathcal{Z}}\frac{\left(\sum_{I\in\mathcal{I}} p_I\nu_I(z)\right)^\alpha}{\left(\sum_{I\in\mathcal{I}} p_I\nu'_I(z)\right)^{\alpha-1}}\mathrm{d}z \\ \le& \frac{1}{\alpha-1}\log\left(\sum_{I\in\mathcal{I}} p_I\int_{\mathcal{Z}}\frac{\nu_I(z)^\alpha}{\nu'_I(z)^{\alpha-1}}\mathrm{d}z\right) \\ =& \frac{1}{\alpha-1}\log\left(\sum_{I\in\mathcal{I}} p_I\exp\big((\alpha-1)D_\alpha(\nu_I\|\nu'_I)\big)\right), \end{aligned}$$

where the inequality is due to a joint convexity argument (see Lemma B.4 in [28]). Similarly, we have

$$D_\alpha(\mathcal{M}(D')\|\mathcal{M}(D)) \le \frac{1}{\alpha-1}\log\left(\sum_{I\in\mathcal{I}} p_I\exp\big((\alpha-1)D_\alpha(\nu'_I\|\nu_I)\big)\right).$$

Let $\mathcal{I}_k := \{I\in\mathcal{I} \mid |D_I\backslash D'_I| = k\}$ denote the set of I such that D_I and D'_I differ by k records, then we have $\sum_{I\in\mathcal{I}_k} p_I = p_k$. Since $\mathcal{A}$ satisfies $(k,\alpha,\tau^*_k(\alpha))$-bounded RGP, both $D_\alpha(\nu'_I\|\nu_I)$ and $D_\alpha(\nu_I\|\nu'_I)$ are upper bounded by $\tau^*_k(\alpha)$ for all $I\in\mathcal{I}_k$. Therefore, it follows that

$$\begin{aligned} &(\alpha-1)\cdot\max\left\{D_\alpha(\mathcal{M}(D)\|\mathcal{M}(D')), D_\alpha(\mathcal{M}(D')\|\mathcal{M}(D))\right\} \\ \le&\log\left(\sum_{k=0}^m\sum_{I\in\mathcal{I}_k} p_I\exp\big((\alpha-1)\cdot\max\left\{D_\alpha(\nu_I\|\nu'_I), D_\alpha(\nu'_I\|\nu_I)\right\}\big)\right) \\ \le&\log\left(\sum_{k=0}^m p_k\exp\Big((\alpha-1)\tau^*_k(\alpha)\Big)\right), \end{aligned}$$

thereby completing the proof for bounded group privacy after substituting the p_k into (10).

Note that in the above analysis, we implicitly assume that the probability density functions μ_k and ν_I are continuous for simplicity. In discrete scenarios, one can readily confirm the validity of the above result by substituting integrals with summations. Therefore, Theorem 3.1 is applicable to both continuous and discrete scenarios.

3.3 Tightness of Our General RGP Bound

Next, we demonstrate that the RGP bound in Theorem 3.1 is asymptotic optimal. Specifically, we first establish a lower bound for the privacy cost of subsampled mechanisms under RGP, and then show that the established lower bound and the upper bound in Theorem 3.1 match up to an additive constant factor. This lower bound is obtained by constructing a pair of binary and one-dimensional m-neighboring datasets D and D', as follows:

$$D = \{0,0,\ldots,0\} \quad\text{and}\quad D' = D\cup\underbrace{\{1,1,\ldots,1\}}_{m\text{ records}}. \tag{11}$$

Let $\mathcal{A}$ be a Gaussian mechanism that takes a binary dataset as the input and output the sum of its records in a differentially private manner, i.e., $\mathcal{A}(D) = \sum_{x\in D} x + \mathcal{N}(0,\sigma^2)$, where $\mathcal{N}(0,\sigma^2)$ denotes the Gaussian noise with mean 0 and variance σ^2. Define $\mathcal{M}(\cdot) := \mathcal{A}\circ\mathsf{Subsample}(\cdot)$ as a subsampled Gaussian mechanism with a sampling rate q, and denote by μ_k the pdf of the distribution $\mathcal{N}(k,\sigma^2)$ for $k\in\mathbb{N}$, then we have $\mathcal{M}(D)\sim\mu_0$ and $\mathcal{M}(D')\sim\sum_{k=0}^m p_k\mu_k$, where $p_k = \binom{m}{k}(1-q)^{m-k}q^k$.

Now suppose that $\mathcal{A}$ satisfies $(k,\alpha,\tau^*_k(\alpha))$-unbounded RGP and $\mathcal{M}$ satisfies $(m,\alpha,\tau_m(\alpha))$-unbounded RGP. This immediately implies that $\tau_m(\alpha)\ge D_\alpha(\mathcal{M}(D')\|\mathcal{M}(D))$. Hence, we obtain:

$$\begin{aligned} \tau_m(\alpha) \ge& \frac{1}{\alpha-1}\log\mathbb{E}_{\mu_0}\left[\left(\frac{\sum_{k=0}^m p_k\mu_k}{\mu_0}\right)^\alpha\right] \\ \ge& \frac{1}{\alpha-1}\log\mathbb{E}_{\mu_0}\left[\sum_{k=0}^m p_k^\alpha\left(\frac{\mu_k}{\mu_0}\right)^\alpha\right] \\ =& \frac{1}{\alpha-1}\log\left(\sum_{k=0}^m p_k^\alpha\mathbb{E}_{\mu_0}\left[\left(\frac{\mu_k}{\mu_0}\right)^\alpha\right]\right) \\ =& \frac{1}{\alpha-1}\log\left(\sum_{k=0}^m p_k^\alpha\exp\Big((\alpha-1)\tau^*_k(\alpha)\Big)\right), \end{aligned} \tag{12}$$

where the last equality holds because a direct calculation verifies that $\mathbb{E}_{\mu_0}[(\mu_k/\mu_0)^\alpha] = \alpha k^2/2\sigma^2$, which attains the RGP guarantee of the Gaussian mechanism $\mathcal{A}$ (see Lemma 4.2). Define

$$\Xi^+_{\alpha,m,q} := \sum_{k=0}^m\binom{m}{k}(1-q)^{m-k}q^k\exp\Big((\alpha-1)\tau^*_k(\alpha)\Big),$$

$$\Xi^-_{\alpha,m,q} := \sum_{k=0}^m\left(\binom{m}{k}(1-q)^{m-k}q^k\right)^\alpha\exp\Big((\alpha-1)\tau^*_k(\alpha)\Big),$$

then by Theorem 3.1 and (12), the upper and lower bounds of the group privacy guarantee of $\mathcal{M}$ can be expressed as $\frac{1}{\alpha-1}\log\Xi^+_{\alpha,m,q}$ and $\frac{1}{\alpha-1}\log\Xi^-_{\alpha,m,q}$, respectively.

We proceed to compare the terms $\Xi^+_{\alpha,m,q}$ and $\Xi^-_{\alpha,m,q}$. Note that

$$\begin{aligned} \sum_{k=0}^m\left(\binom{m}{k}(1-q)^{m-k}q^k\right)^\alpha \ge& \left(\binom{m}{0}(1-q)^m q^0\right)^\alpha \\ =&(1-q)^{m\alpha} \ge 1-qm\alpha, \end{aligned} \tag{13}$$

where the last inequality follows from the Bernoulli's inequality. Define $c_\alpha := \max\{(\alpha-1)\tau^*_k(\alpha)\}_{k=0}^m$. With an appropriate setting for the injected noise to satisfy group privacy (e.g., $\sigma=\Theta(m)$ in Gaussian mechanisms, elaborate later in Section 4.1), c_α becomes a constant depending solely on α. Thus, we derive:

$$\begin{aligned} \Xi^+_{\alpha,m,q} - \Xi^-_{\alpha,m,q} =& \sum_{k=0}^m\Big(p_k - p_k^\alpha\Big)e^{((\alpha-1)\tau^*_k(\alpha))} \\ \le& e^{c_\alpha}\left(\sum_{k=0}^m p_k - \sum_{k=0}^m p_k^\alpha\right) = e^{c_\alpha}\left(1-\sum_{k=0}^m p_k^\alpha\right) \le e^{c_\alpha}qm\alpha, \end{aligned}$$

where the last inequality follows from (13). By applying the mean value theorem to the logarithm function, we can verify that the upper and lower bounds, i.e., $\frac{1}{\alpha-1}\log\Xi^+_{\alpha,m,q}$ and $\frac{1}{\alpha-1}\log\Xi^-_{\alpha,m,q}$, match up to an additive factor $O(e^{c_\alpha}qm\alpha)$. Setting $qm=O(1)$ and treating terms influenced by α as constants implies that the upper and lower bounds match up to an additive constant factor, implying the asymptotic tightness of our bound.

Remark. As we show soon in Section 4.1, the RGP upper and lower bounds discussed above for the DP-SGD algorithm match up to an additive factor of $O(\alpha e^{\alpha^2})$, provided that the noise scale $\sigma=\Theta(m)$

and $qm = O(1)$. Further, although the above description uses the subsampled Gaussian mechanism as an example, the analysis of the lower bound can be adapted for other mechanisms, which is detailed in Appendix D of the full version [28].

While asymptotic results provide valuable insights on the optimality of our general RGP bound, for practical applications, we also need concrete bounds for calibrating the noise levels. In the next section, we further derive closed-form group privacy bounds for a range of widely-used subsampled mechanisms.

4 APPLICATIONS

This section presents closed-form group privacy bounds for various privacy-preserving subsampled mechanisms. In what follows, we denote by f the algorithm intended for privatization under RGP.

We first present a lemma which is instrumental in deriving tight RGP bounds for non-subsampled mechanisms, i.e., the term $\tau_k^*(\alpha)$ in Theorem 3.1. This lemma serves as a useful tool for deriving closed-form RGP guarantees for subsampled Gaussian, Laplace, and Skellam mechanisms, detailed in the following subsections.

Lemma 4.1. *Let $f : \mathcal{D} \mapsto \mathbb{R}^d$ be a function that maps a dataset to a d-dimensional vector and let $\|\cdot\|$ be an arbitrary norm. If $\|f(D) - f(\hat{D})\| \leq C$ holds for any pair of bounded (resp. unbounded) 1-neighboring datasets D and $\hat{D}$ that differ by one record, then for all $k \in \mathbb{N}$, it holds that $\|f(D) - f(D')\| \leq kC$ for any pair of bounded (resp. unbounded) k-neighboring datasets D and D'.*

Proof. The lemma is established by applying the triangle inequality of norms and the principle of mathematical induction. □

4.1 Subsampled Gaussian Mechanism

We now establish the closed-form RGP bound for the subsampled Gaussian mechanism, which is a fundamental component in many popular privacy-preserving applications, most notably DP-SGD [1]. Let $f : \mathcal{D} \mapsto \mathbb{R}^d$ denote the algorithm for which $\|f(D) - f(\hat{D})\|_2 \leq C$ holds for all pairs of datasets differing by one record. The Gaussian mechanism that privatizes the algorithm f is defined as follows:

$$\mathcal{A}(D) = f(D) + \mathcal{N}(\mathbf{0}, C^2\sigma^2\mathbb{I}^d),$$

where $\mathcal{N}(\mathbf{0}, C^2\sigma^2\mathbb{I}^d)$ represents the spherical d-dimensional Gaussian noise with per-coordinate variance $C^2\sigma^2$.

For any pair of k-neighboring dataset D and D', the Rényi divergence between $\mathcal{A}(D)$ and $\mathcal{A}(D')$ can be upper bounded as:

$$\begin{aligned}
&D_\alpha\Big(\mathcal{N}\Big(f(D), C^2\sigma^2\mathbb{I}^d\Big)\Big\|\mathcal{N}\Big(f(D'), C^2\sigma^2\mathbb{I}^d\Big)\Big)\\
&\overset{(a)}{=} D_\alpha\Big(\mathcal{N}\Big(f(D) - f(D'), C^2\sigma^2\mathbb{I}^d\Big)\Big\|\mathcal{N}\Big(\mathbf{0}, C^2\sigma^2\mathbb{I}^d\Big)\Big)\\
&\overset{(b)}{\leq} \sup_{\|\mathbf{v}\|_2 \leq kC} D_\alpha\Big(\mathcal{N}(\mathbf{v}, C^2\sigma^2\mathbb{I}^d)\Big\|\mathcal{N}(\mathbf{0}, C^2\sigma^2\mathbb{I}^d)\Big)\\
&\overset{(c)}{=} \sup_{\|\mathbf{v}\|_2 \leq kC} \sum_{i=1}^{d} D_\alpha\Big(\mathcal{N}(\mathbf{v}[i], C^2\sigma^2)\Big\|\mathcal{N}(\mathbf{0}, C^2\sigma^2)\Big)\\
&\overset{(d)}{=} \sup_{\|\mathbf{v}\|_2 \leq kC} \sum_{i=1}^{d} \frac{\alpha\mathbf{v}[i]^2}{2C^2\sigma^2} = \sup_{\|\mathbf{v}\|_2 \leq kC} \frac{\alpha\|\mathbf{v}\|_2^2}{2C^2\sigma^2} = \frac{\alpha k^2}{2\sigma^2}, \qquad (14)
\end{aligned}$$

where (a) follows from the invariance of Rényi divergence under invertible transformations, which is a variation of the more general data processing inequality [50]; (b) is derived from Lemma 4.1; (c) follows from the additivity of Rényi divergence (see Lemma B.3 in [28]); (d) follows from the closed-form Rényi divergence between Gaussian distributions (see Lemma B.5 in [28]). Here, $\mathbf{v}[i]$ denotes the i-th element of vector $\mathbf{v}$. It is worth noting that the upper bound in (14), i.e., $\alpha k^2/2\sigma^2$, can be attained by specific f and D, affirming the tightness of our analysis.

Accordingly, we establish the RGP guarantee for the Gaussian mechanism as follows:

Lemma 4.2. *Let $\mathcal{A}$ be a Gaussian mechanism defined as above, it holds that $\mathcal{A}$ satisfies $(k, \alpha, \tau = \frac{\alpha k^2}{2\sigma^2})$-RGP for all $k \in \mathbb{N}$.*

Comparing the above result with the generic RDP-to-RGP conversion in Lemma 2.3, observe that the former applies to any group size, whereas the latter is limited to the case where the group size is a power of two, i.e., 2^c for $c \in \mathbb{N}$. Further, Lemma 4.2 above leads to the same value of α in both the RDP (i.e., by setting $k = 1$) and RGP ($k > 1$), where as in Lemma 2.3, the converted RGP uses a lower $\alpha' = \frac{\alpha}{2^c}$. If, in Lemma 4.2, we also aim to satisfy RGP with this smaller α', with $k = 2^c$, we would have $\tau = \frac{\alpha' k^2}{2\sigma^2} = \frac{\alpha k}{2\sigma^2}$, which grows linearly with the group size k, rather than with $k^{1.58}$ as discussed in Section 2.2. Hence, the above lemma provides a more refined privacy analysis compared to the naive RDP-to-RGP conversion approach.

Next, we extend our refined, closed-form privacy cost analysis to the subsampled Gaussian mechanism. Combining Lemma 4.2 with Theorem 3.1, we arrive at the following theorem.

Theorem 4.1. *Let $\mathcal{M} := \mathcal{A} \circ$ Subsample denote the subsampled Gaussian mechanism with sampling rate $q \in (0, 1)$, where $\mathcal{A}$ is the Gaussian mechanism defined above. Then, $\mathcal{M}$ satisfies $(m, \alpha, \tau_m(\alpha))$-RGP, with*

$$\tau_m(\alpha) = \frac{1}{\alpha - 1} \log\left(\sum_{k=0}^{m} \binom{m}{k} (1-q)^{m-k} q^k \exp\left(\frac{(\alpha-1)\alpha k^2}{2\sigma^2}\right)\right).$$

The following theorem shows that with respect to the subsampled Gaussian mechanism, the above result achieves an asymptotically improved privacy guarantee compared to baseline solution of RDP-to-RGP conversion through Lemma 2.3.

Theorem 4.2. *Consider a subsampled Gaussian mechanism $\mathcal{M}$ with a sampling rate $q \in (0, 1)$ and variance σ^2. Let $\alpha > 1$ be any integer. We denote by $(m, \alpha, \tau'_m(\alpha))$ and $(m, \alpha, \tau_m(\alpha))$ the unbounded RGP guarantees of $\mathcal{M}$ derived using Lemma 2.3 and Theorem 4.1, respectively. If $\sigma = \Theta(m)$ and $qm > 1$, our bound $\tau_m(\alpha)$ saves a multiplicative factor of $\Theta(m^{\log_2 1.5})$ compared to $\tau'_m(\alpha)$.*

Proof sketch. The complete proof is deferred to Appendix C.1 of the full version [28]. Here we present the proof sketch as follows.

Both the RGP bounds derived by the baseline solution and our solution can be expressed in a binomial expression-like form that includes exponential terms. These terms are rather complicated to compare directly. To address this challenge, we approximate these exponential terms using polynomials via Taylor's theorem. This approximation allows the RGP bounds to be asymptotically expressed as the summation of high moments of binomial distributions, yielding simpler and clearer bounds. Accordingly, we derive asymptotic improvements based on these refined bounds. □

Remark. The above theorem indicates that our RGP bound significantly enhances the privacy guarantee for the DP-SGD algorithm, on the condition that $\sigma = \Theta(m)$ and $qm > 1$. Note that these two conditions align with the conditions in the remark toward the end of Section 3.3, signifying that our bound is provably tight for DP-SGD while saving a $\Theta(m^{\log_2 1.5}) \approx \Theta(m^{0.58})$ multiplicative factor[4] in terms of RGP guarantee compared to the baseline method.

4.2 Subsampled Laplace Mechanism

The Laplace mechanism [20] is among the most widely used mechanisms for achieving DP in numerous practical applications [8, 17, 21, 39, 58]. Let $f : \mathcal{D} \mapsto \mathbb{R}^d$ denote the algorithm for which $\|f(D) - f(\hat{D})\|_1 \leq C$ is satisfied for all pairs of neighboring datasets that differ by one record. The Laplace mechanism is defined as

$$\mathcal{A}(D) = f(D) + \mathsf{Lap}(\mathbf{0}, Cb\mathbb{I}^d),$$

where $\mathsf{Lap}(\mathbf{0}, Cb\mathbb{I}^d)$ denotes the d-dimensional Laplace noise with per-coordinate scale factor Cb, i.e., each coordinate is independently drawn from the one-dimensional Laplace distribution $\mathsf{Lap}(0, Cb)$.

We can upper bound the Rényi divergence between $\mathcal{A}(D)$ and $\mathcal{A}(D')$ for any pair of k-neighboring datasets D and D' as follows:

$$\begin{aligned}
&D_\alpha(\mathsf{Lap}(f(D), Cb\mathbb{I}^d) \| \mathsf{Lap}(f(D'), Cb\mathbb{I}^d)) \\
&\overset{(a)}{=} D_\alpha\Big(\mathsf{Lap}(f(D) - f(D'), Cb\mathbb{I}^d) \Big\| \mathsf{Lap}(\mathbf{0}, Cb\mathbb{I}^d)\Big) \\
&\overset{(b)}{\leq} \sup_{\|\mathbf{v}\|_1 \leq kC} \sum_{i=1}^{d} D_\alpha(\mathsf{Lap}(\mathbf{v}[i], Cb) \| \mathsf{Lap}(0, Cb)) \\
&\overset{(c)}{=} \sup_{\|\mathbf{v}\|_1 \leq kC} \frac{1}{\alpha - 1} \sum_{i=1}^{d} \log\left\{ \frac{\alpha}{2\alpha - 1} \exp\left(\frac{(\alpha - 1)\mathbf{v}[i]}{Cb} \right) \right. \\
&\qquad\qquad \left. + \frac{\alpha - 1}{2\alpha - 1} \exp\left(\frac{-\alpha \mathbf{v}[i]}{Cb} \right) \right\}, \qquad (15)
\end{aligned}$$

where (a) follows from the invariance of Rényi divergence under invertible transformations; (b) follows from Lemma 4.1 and the additivity of Rényi divergence (see Lemma B.3 in the full version [28]); (c) follows from the closed-form Rényi divergence between Laplace distributions (see Lemma B.6 in [28]).

The subsequent task is to derive a closed-form upper bound of (15). Unlike the Gaussian mechanism, the expression in (15) is rather complicated, and thus identifying its maximum is not straightforward. To address this, we formulate the following constrained optimization problem, which is equivalent to determining the maximum of (15):

$$\begin{aligned}
\text{maximize}_{\{x_i\}} \quad & \sum_{i=1}^{d} \log\left\{ \frac{\alpha}{2\alpha - 1} \exp\left(\frac{(\alpha - 1)x_i}{Cb} \right) \right. \\
& \left. + \frac{\alpha - 1}{2\alpha - 1} \exp\left(\frac{-\alpha x_i}{Cb} \right) \right\} \\
\text{subject to} \quad & \sum_{i=1}^{d} |x_i| \leq kC. \qquad (16)
\end{aligned}$$

Next, we introduce the following lemmas, which are crucial tools for solving the above constraint optimization problem.

Lemma 4.3. *The function* $f(x) = \log\left(c_1 e^{\beta_1 x} + c_2 e^{-\beta_2 x}\right)$ *is convex for* $c_1, c_2, \beta_1, \beta_2 \in (0, \infty)$.

Proof. See Appendix C.2 of the full version [28]. □

Lemma 4.4 (Bauer's Maximum Principle [9])**.** *A maximum of a convex function over a closed and bounded convex set is achieved at an extreme point.*

We are now ready to solve the optimization problem (15). Note that for all $\alpha > 1$, we have $\frac{\alpha}{2\alpha-1} > 0$, $\frac{\alpha-1}{2\alpha-1} > 0$, $\frac{\alpha-1}{Cb} > 0$, and $\frac{\alpha}{Cb} > 0$. Thus, by Lemma 4.3, the objective function of (16) is convex. Moreover, the domain of this objective function is a convex polyhedron (an L_1-ball), with vertices constituting its extreme points. Consequently, by Lemma 4.4, the maximum of (15) is attained at $\mathbf{v} = kC \cdot \mathbf{e}$ for some vector $\mathbf{e}$ in the standard basis. Denote by $\mathbf{e}_i$ the i-th element of the vector $\mathbf{e}$, then we have

$$\begin{aligned}
(15) \leq & \frac{1}{\alpha - 1} \sum_{i=1}^{d} \log\left\{ \frac{\alpha}{2\alpha - 1} \exp\left(\frac{(\alpha - 1)kC\mathbf{e}_i}{Cb} \right) \right. \\
& \left. + \frac{\alpha - 1}{2\alpha - 1} \exp\left(\frac{-\alpha kC\mathbf{e}_i}{Cb} \right) \right\} \\
= & \frac{1}{\alpha - 1} \log\left\{ \frac{\alpha}{2\alpha - 1} \exp\left(\frac{(\alpha - 1)k}{b} \right) + \frac{\alpha - 1}{2\alpha - 1} \exp\left(\frac{-\alpha k}{b} \right) \right\}. \qquad (17)
\end{aligned}$$

Substituting (17) into Theorem 3.1 establishes the following group privacy guarantee for the subsampled Laplace mechanism:

Theorem 4.3. *Let* $\mathcal{M} \coloneqq \mathcal{A} \circ \mathsf{Subsample}$ *be a subsampled Laplace mechanism with sampling rate* $q \in (0, 1)$*, where* $\mathcal{A}$ *is the Laplace mechanism defined above. Then,* $\mathcal{M}$ *satisfies* $(m, \alpha, \tau_m(\alpha))$*-RGP, where*

$$\tau_m(\alpha) = \frac{1}{\alpha - 1} \log\left(\sum_{k=0}^{m} \binom{m}{k} (1 - q)^{m-k} q^k \Phi_k^{\mathsf{Lap}}(\alpha) \right),$$

with

$$\Phi_k^{\mathsf{Lap}}(\alpha) = \frac{\alpha}{2\alpha - 1} \exp\left(\frac{(\alpha - 1)k}{b} \right) + \frac{\alpha - 1}{2\alpha - 1} \exp\left(\frac{-\alpha k}{b} \right).$$

Remark. The above result is not only interesting for RGP, but also for the RDP setting. Before this work, the best known RDP guarantee to our knowledge for d-dimensional Laplace mechanism (i.e., Theorem 11 in [13]) accumulates privacy cost across all d-dimensions, i.e., by applying privacy composition d times, which incurs a multiplicative factor of d. In contrast, our analysis reduces the problem to a simpler, one-dimensional case by carefully examining the conditions leading to the worst-case scenario. This approach allows us to derive a privacy guarantee without using composition, thus achieving a significant improvement.

Meanwhile, it is also important to note that the fact that our analysis bypasses privacy composition does not imply a violation of the well-established $\Omega(\sqrt{d})$ error bound of $\mathcal{A}$ from [16]. The

[4] Based on our analysis in Appendix C.1 of the full version [28], we can derive that when the Gaussian variance satisfies $\sigma = \Theta(m)$, our RGP guarantee is upper bounded by $\alpha q^2/2$ as m becomes large, while the baseline RGP guarantee is lower bounded by $m^{\log_2 1.5} \alpha q^2/4 \approx m^{0.58} \alpha q^2/4$. Therefore, a more accurate bound on the improvement is $m^{0.58}/2$. In Theorem 4.2, we use the big-oh notation to demonstrate the asymptotic improvement of our result, simplifying the expression by ignoring constant factors.

Laplace noise is proportional to the L_1 norm of the output, and thus, the error associated with the Laplace mechanism still (implicitly) depends on the output dimension d.

4.3 Subsampled Skellam Mechanism

The Skellam mechanism [2, 7] applies discrete noise following a symmetric Skellam distribution to achieve DP. This mechanism is particularly useful in federated learning frameworks operating with multi-party computation (MPC) protocols [10, 54, 59], since these protocols use modular arithmetic as the fundamental cryptographic primitive [57], hence limiting their computation in finite fields. As a result, common mechanisms that typically inject real-valued noise, including the Gaussian and the Laplace mechanisms discussed in previous subsections, are inherently incompatible with these protocols. We refer interested readers to [2, 7] for more details.

Let $\mathsf{Sk}(z, \mu)$ denote the one-dimensional symmetric Skellam distribution with mean $z \in \mathbb{Z}$ and variance 2μ. Denote by $f : \mathcal{D} \mapsto \mathbb{Z}^d$ the algorithm that maps an input dataset to an integer-valued vector such that $\|f(D) - f(\hat{D})\|_1 \leq C$ for pair of datasets D and $\hat{D}$ that differ by one record. The Skellam mechanism is defined as follows:

$$\mathcal{A}(D) = f(D) + \mathsf{Sk}(\mathbf{0}, C^2\mu\mathbb{I}^d),$$

where $\mathsf{Sk}(\mathbf{0}, C^2\mu\mathbb{I}^d)$ denotes the multi-dimensional Skellam distribution with each coordinate distributed independently as $\mathsf{Sk}(0, C^2\mu)$.

The Rényi divergence $D_\alpha(\mathcal{A}(D)\|\mathcal{A}(D'))$ for any pair of k-neighboring datasets D and D' can be bounded as:

$$\begin{aligned}
&D_\alpha(\mathsf{Sk}(f(D), C^2\mu\mathbb{I}^d)\|\mathsf{Sk}(f(D'), C^2\mu\mathbb{I}^d)) \\
&\overset{(a)}{=} D_\alpha(\mathsf{Sk}(f(D) - f(D'), C^2\mu\mathbb{I}^d)\|\mathsf{Sk}(\mathbf{0}, C^2\mu\mathbb{I}^d)) \\
&\overset{(b)}{\leq} \sup_{\|\mathbf{v}\|_1 \leq kC} \sum_{i=1}^{d} \left(\frac{\alpha \mathbf{v}[i]^2}{2C^2\mu} + \min\left\{ \frac{(2\alpha-1)\mathbf{v}[i]^2 + 6|\mathbf{v}[i]|}{4C^4\mu^2}, \frac{3|\mathbf{v}[i]|}{2C^2\mu} \right\} \right) \\
&\overset{(c)}{\leq} \sup_{\|\mathbf{v}\|_1 \leq kC} \left(\frac{\alpha\|\mathbf{v}\|_2^2}{2C^2\mu} + \min\left\{ \frac{(2\alpha-1)\|\mathbf{v}\|_2^2 + 6\|\mathbf{v}\|_1}{4C^4\mu^2}, \frac{3\|\mathbf{v}\|_1}{2C^2\mu} \right\} \right) \\
&\overset{(d)}{\leq} \frac{\alpha k^2}{2\mu} + \min\left\{ \frac{(2\alpha-1)k^2C + 6k}{4C^3\mu^2}, \frac{3k}{2C\mu} \right\}, \qquad (18)
\end{aligned}$$

where (a) follows from the invariance of Rényi divergence under invertible transformations; (b) is derived using Lemma 4.1, the closed-form Rényi divergence between symmetric Skellam distributions (see Lemma B.7 in [28]), and the additivity of Rényi divergence (see Lemma B.3 in [28]); (c) follows from the definitions of L_1 and L_2 norms; (d) follows from the fact that $\|\mathbf{v}\|_1 \leq C$ implies $\|\mathbf{v}\|_2 \leq C$. Combining (18) with Theorem 3.1 leads to the following group privacy guarantee for subsampled Skellam mechanisms:

Theorem 4.4. *Let $\mathcal{M}$ be the subsampled Skellam mechanism as defined above. $\mathcal{M}$ then satisfies $(m, \alpha, \tau_m(\alpha))$-RGP, where*

$$\tau_m(\alpha) = \frac{1}{\alpha - 1} \log\left(\sum_{k=0}^{m} \binom{m}{k} (1-q)^{m-k} q^k \exp\left((\alpha - 1)\Phi_k^{\mathsf{Sk}}(\alpha) \right) \right),$$

with

$$\Phi_k^{\mathsf{Sk}}(\alpha) = \frac{\alpha k^2}{2\mu} + \min\left\{ \frac{(2\alpha-1)k^2C + 6k}{4C^3\mu^2}, \frac{3k}{2C\mu} \right\}.$$

4.4 Subsampled Randomized Response

The randomized response (RR) mechanism [52] allow for the collection of statistical information about sensitive datasets while ensuring a DP guarantee. Given a predicate function $f : \mathcal{D} \mapsto \{0, 1\}$, the randomized response mechanism $\mathcal{A}$ that privatizes f with privacy parameter $p \in (0.5, 1)$ is defined as follows:

$$\mathcal{A}(D) = \begin{cases} f(D), & \text{w.p.} \quad p \\ 1 - f(D), & \text{w.p.} \quad 1 - p. \end{cases} \qquad (19)$$

Note that the RR mechanism provides a local DP guarantee [22], ensuring that the outputs of RR mechanisms are indistinguishable for any pair of datasets that differ by any number of records. Therefore, the RR mechanism satisfies $(m, \alpha, \tau_1^*(\alpha))$-RGP for any $m \in \mathbb{N}$. According to Lemma B.8 in the full version [28], it holds that

$$\tau_1^*(\alpha) = \frac{1}{\alpha - 1} \log\left(\frac{p^\alpha}{(1-p)^{\alpha-1}} + \frac{(1-p)^\alpha}{p^{\alpha-1}} \right). \qquad (20)$$

In addition, a direct calculation yields $\tau_0^*(\alpha) = 0$. Consequently, we can establish the following group privacy guarantee for the subsampled RR mechanism:

Theorem 4.5. *The subsampled RR mechanism with sampling rate q and privacy parameter p satisfies $(m, \alpha, \tau_m(\alpha))$-group RDP, where*

$$\tau_m(\alpha) = \frac{1}{\alpha - 1} \log\left((1-q)^m + \left(1 - (1-q)^m\right) \cdot \Phi^{\mathsf{RR}}(\alpha) \right),$$

where

$$\Phi^{\mathsf{RR}}(\alpha) = \frac{p^\alpha}{(1-p)^{\alpha-1}} + \frac{(1-p)^\alpha}{p^{\alpha-1}}.$$

The following lemma shows that our group privacy bound for the RR mechanism achieves a provable improvement over naive bound $\tau_1^*(\alpha)$ defined in Eq. (20).

Lemma 4.5. *For the RR mechanism, it holds that $\lim_{m\to\infty} \tau_m(\alpha) = \tau_1^*(\alpha)$. In addition, we have $\tau_m(\alpha) < \tau_1^*(\alpha)$ for all $m \in \mathbb{N}$.*

Proof. See Appendix C.3 of the full version [28]. □

5 EXPERIMENTS

This section presents comparative experiments between our proposed RGP bound and the existing method with generic DP-to-GP conversion. We start by comparing the required noise levels for achieving a given (m, α, τ)-RGP guarantee through different RGP bounds in Section 5.1. Then, we evaluate the practicability of our RGP solution using the DP-SGD algorithm for neural network model training in Section 5.2. Specifically, the RGP bounds examined in our experiments are listed as follows:

(1) The baseline method, which is derived by first obtaining the RDP guarantee of the subsampled mechanism according to the RDP bound in [61], then converting the RDP guarantee to the RGP guarantee using Lemma 2.3.
(2) Our closed-form RGP bounds in Section 4.

In addition, to demonstrate the tightness of our RGP bounds, we compare the noise levels required by our bounds with the noise levels calibrated by the analytical lower bounds across different subsampled mechanisms. The details of the analytical lower bounds can be found in Appendix D of the full version [28].

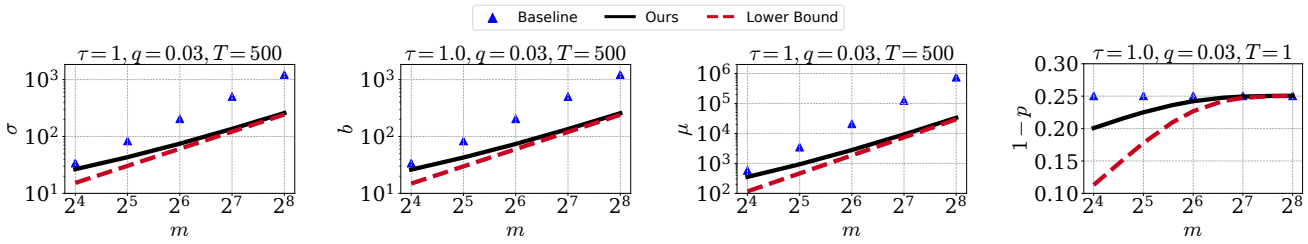

(a) Subsampled Gaussian (b) Subsampled Laplace (c) Subsampled Skellam (d) Subsampled RR

Figure 2: Varying the group size m.

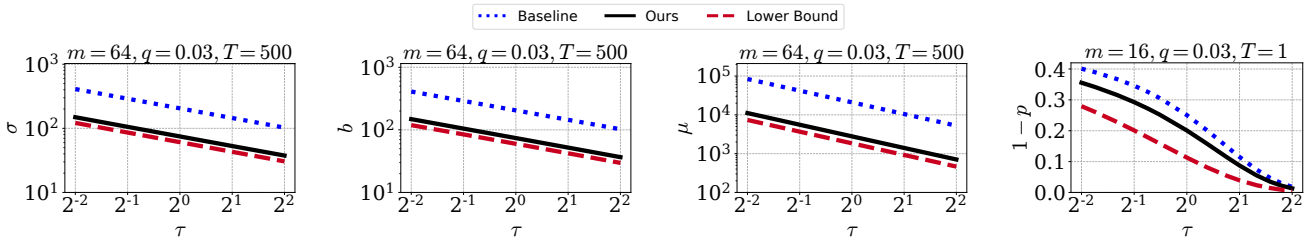

(a) Subsampled Gaussian (b) Subsampled Laplace (c) Subsampled Skellam (d) Subsampled RR

Figure 3: Varying the privacy parameter τ.

Note that the closed-form RGP guarantees in Section 4 apply to both unbounded and bounded RGP notions. From the perspective of algorithmic implementation, ensuring unbounded (or bounded) RGP necessitates bounding the difference $\|f(D) - f(\hat{D})\|$ for all pairs of unbounded (or bounded) neighbor datasets D and $\hat{D}$ by a constant C. In light of this, we specify the value of C but do not differentiate between unbounded and bounded cases. The comparative results between unbounded and bounded RGP bounds are identical in the experiments.

5.1 Comparison of Noise Levels

In this subsection, we conduct a numerical comparison of RGP bounds across different configurations of the Gaussian, Laplace, Skellam, and RR mechanisms. Specifically, for a given (m, α, τ)-RGP guarantee, we calibrate the required noise parameters (σ for Gaussian, b for Laplace, μ for Skellam, and p for the RR mechanism) to achieve this guarantee through our proposed analysis, the baseline solution, and the theoretical lower bound. In the experiments, we set the norm bound C to 1 for the Skellam mechanism; for Gaussian and Laplace mechanisms, the norm bound is not specified as it is canceled out in the final expression of the RGP guarantee.

Varying the group size m. We first vary the group size from 16 to 256 while keeping other parameters fixed, and then compare the required noise levels derived from our analysis, the baseline method, and the theoretical lower bound. Specifically, we fix the privacy guarantee to $(m, \alpha = 4, \tau = 1)$-RGP, which can be converted to $(m, \epsilon = 4.1, \delta = 10^{-5})$-GP via Lemma 2.2. The results are depicted in Figure 2. Details on the sampling rate q, and the number of iterations T are provided in each subfigure. Note that in the RR mechanism, a lower p value indicates larger noise. Therefore, we compare the value of $1 - p$ for the RR mechanism to facilitate a better comparison of noise levels.

Our main observations from the experimental results are threefold: (i) for Gaussian, Laplace, and Skellam mechanisms, our bound is not only tight but also clearly superior to the baseline; (ii) as m increases, the gap between our bound and the baseline method grows for the Gaussian mechanism, aligning with the theoretical insights from Theorem 4.2; (iii) for the RR mechanism, while our bound consistently outperforms the baseline, there is a gap between our bound and the theoretical lower bound at small group sizes, suggesting room for further improvement in RGP guarantees.

Varying the privacy parameter τ. We vary the RGP privacy parameter τ by {0.25, 0.5, 1, 2, 4} and calibrate the noise levels to achieve $(m, \alpha = 4, \tau)$-RGP guarantee. These RGP guarantees can be converted (i.e., using Lemma 2.2) to $(\epsilon, \delta = 10^{-5})$-GP guarantees, with the corresponding ϵ values being 3.4, 3.6, 4.1, 5.2, and 7.1, respectively. The experimental results are presented in Figure 3. From these results, we observe that the required noise levels calibrated by our bounds are consistently and significantly smaller than those required by the baseline method. Furthermore, for Gaussian, Laplace, and Skellam mechanisms, our privacy analysis results tightly match the corresponding theoretical lower bounds.

We also compare the RGP bounds by varying the sampling rate and the number of iterations. The experimental results are reported in Appendix F of the full version [28].

Summary of results. All above experimental results validate that our RGP bound is consistently tight across a wide range of mechanism configurations. Furthermore, these results show over an order of magnitude improvement of our privacy bound compared to the

baseline solution, which enables data privacy practitioners to develop practical mechanisms that offer meaningful RGP guarantees while preserving desirable utility levels. In addition, results of the lower bound indicate that ensuring RGP inevitably requires the injection of noise proportional to the group size, which can significantly impair the utility of the mechanism. Therefore, an interesting direction for future research is to explore new methods to relax and redefine the notion of group privacy, which can help design privacy-preserving mechanisms that offer a more favorable balance between group privacy protection and result utility.

5.2 Evaluations on DP-SGD

In this subsection, we conduct experiments on model training via the DP-SGD algorithm [1] with different privacy guarantees to validate the practicability and superiority of our RGP bound. It is important to note that the DP-SGD algorithm essentially operates as a type of subsampled Gaussian mechanism [43]. Therefore, given specific privacy parameters, we can calibrate the required Gaussian noise variance based on our closed-form RGP bounds for subsampled Gaussian mechanisms, as presented in Theorem 4.1. Our experiments involve comparisons between two different implementations of the DP-SGD algorithm: one implementation calibrates the noise to achieve group privacy according to Theorem 4.1, and the other calibrates noise based on the baseline solution. Note that DP-SGD offers (m, ϵ, δ)-unbounded GP guarantees [1]. For consistency, we also focus on the (m, ϵ, δ)-unbounded GP in the following experiments.

Setup. We conduct experiments on three benchmark image classification datasets: MNIST [34], Fashion-MNIST [55], and CIFAR-10 [33]. Note that in many practical scenarios, image classification tasks require formal group privacy guarantees. For example, it is necessary to prevent the inference of gender or race ratios within a specific group in human face datasets [38]. Similarly, it is essential to safeguard information about genetic diseases within groups, such as families, in medical image datasets.

Following the state-of-the-art differentially private learning solution [48], we use logistic regression (LR) and convolutional neural networks (CNN) paired with Scattering Networks (SN) in our experiments. We defer the details of the experimental setup to Appendix E. In all experiments, we fix the privacy parameter δ to 10^{-5}.

To achieve an $(m, \epsilon, \delta = 10^{-5})$-GP guarantee, we employ a binary search method to determine the required Gaussian noise scale σ. For each noise scale σ, we initially obtain (m, α, τ)-RGP guarantees for each α in the set $\{2, 3, \ldots, 100\}$ based on our closed-form RGP bound for subsampled Gaussian mechanism in Theorem 4.1. Subsequently, we calculate the corresponding privacy parameters ϵ for each α by Lemma 2.2. The minimum ϵ identified in this process is then selected as the GP guarantee.

Varying the group size m. In this experiment, we set the group privacy parameter ϵ to 4 and vary the group size m by $\{8, 16, 24, 32, 40, 48, 56, 64\}$ to show the impact of different group sizes on the model utility. It is important to note that the baseline solution in Lemma 2.3 can only ensure RGP for group sizes that are powers of two. Therefore, we vary the group size m to $\{8, 16, 32, 64\}$ for the baseline solution.

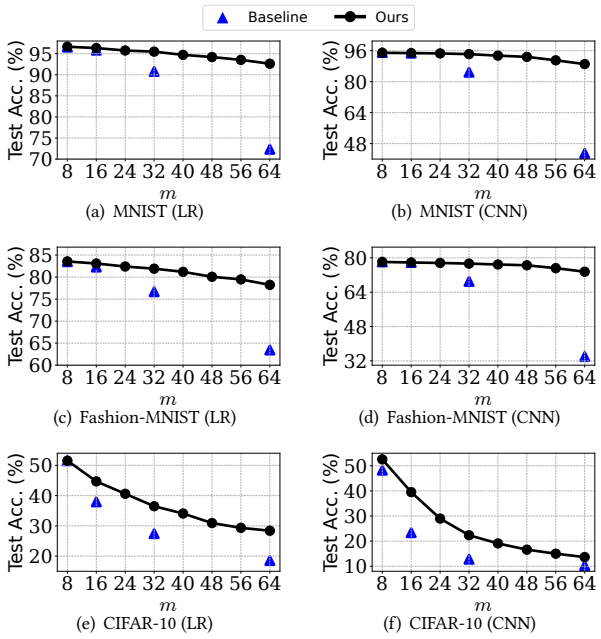

Figure 4: DP-SGD with $(m, 4, 10^{-5})$-GP guarantee.

Figure 4 presents the model test accuracy under different group sizes. We observe that for small group sizes, e.g., $m \leq 16$, our solution performs similarly to the baseline solution. However, for larger group sizes such as $m \geq 32$, our solution consistently and significantly outperforms the baseline solution, while maintaining a desirable model utility. In addition, for all experimental results except for CNN on the CIFAR-10 dataset, we see that the model accuracy gap between our solution and the baseline grows as m increases, confirming our theoretical insights from Theorem 4.2 and aligning with the numerical results in Figure 2(a). For the experimental result of CNN on CIFAR-10 in Figure 4(f), with a smaller group size, e.g., when $m \leq 16$, the model accuracy gap between the proposed and the baseline solutions grows with m. When m is relatively large (e.g., over 32), however, the gap no longer expands since the noise levels calibrated by the baseline solution become too large for the deep CNN model to converge, leading to model accuracy close to random guessing (i.e., $\approx 10\%$).

Varying the privacy parameter ϵ. Figure 5 reports the model test accuracy for a fixed group size of $m = 32$ across different privacy parameters ϵ. We observe that our solution consistently and significantly outperforms the baseline solution for all settings of ϵ. Specifically, while achieving the same level of group privacy guarantee, our solution dramatically reduces the noise injected into the model weights compared to the baseline method, thus producing models with higher utility, which are more suitable for their target real-world applications.

6 RELATED WORK

As a natural extension of DP, the notion of group privacy provides formal protection for the aggregate information of a group of individuals and has been widely recognized and explored in recent

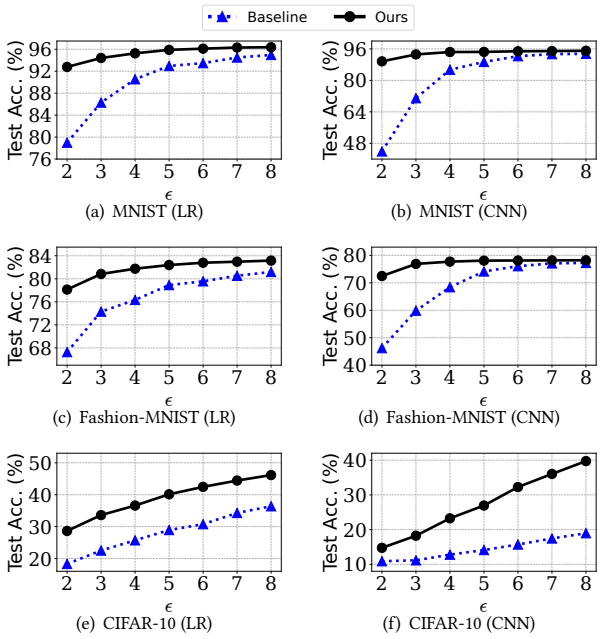

Figure 5: DP-SGD with $(32, \epsilon, 10^{-5})$-GP guarantee.

years [21, 41, 47, 49]. Previous studies [21, 42, 49] have developed general methodologies to convert DP guarantees to GP guarantees. Nonetheless, as explained in Sections 1 and 2.2, these methods treat DP mechanisms as black boxes to ensure general applicability, resulting in overly conservative results. This limitation has been a driving motivation behind this paper.

We note that a concurrent work [46] also studies RGP bounds for subsampled mechanisms. The main differences between our work and this concurrent study are as follows. First, the RGP bound in [46] involves integrals on a rather complicated function, in which both the numerator and denominator are higher powers of the PDFs of mixture distributions. The number of mixture components — and thus the complexity of the integral function — increases with the group size, leading to numerical instability and inaccuracies for large groups. In contrast, our work provides closed-form RGP bounds that can be efficiently computed with numerical tools. Second, the RGP bound in [46] requires non-trivial and *mechanism-specific* theoretical privacy analysis for each mechanism. This complexity makes it challenging to apply to mechanisms with intricate noise distributions, such as the Skellam mechanism. On the contrary, our bound is concise, requiring only the derivation of RGP guarantees for the corresponding non-subsampled mechanism. As a result, it can be easily applied to derive closed-form RGP bounds for various subsampled mechanisms, as demonstrated in Section 4. Finally, we established analytical lower bounds to empirically verify the tightness of our bounds and provided experimental evaluations on logistic regression and deep neural networks using benchmark datasets, demonstrating the practical significance of our contributions. In contrast, the concurrent work mainly focuses on theoretical analysis and does not provide any empirical evaluations.

In addition to group privacy, there have been several studies on the protection of user-level privacy in recent years. Particularly, user-level DP [3, 11, 23, 24, 35, 37, 40] has garnered considerable attention. Various practical algorithms ensuring user-level DP have been proposed in [3, 23, 40]. Meanwhile, some studies [11, 24, 37] primarily focus on the theoretical aspects. Notable among these are recent works [11, 24], which have delved into methods of converting DP guarantees into user-level DP guarantees. It is crucial to note, however, that these studies rely on asymptotic analysis with additional assumptions, such as i.i.d. distribution of data records. Consequently, while these works contribute valuable theoretical insights, they fall short of providing precise, closed-form privacy guarantees that are essential for designing practical privacy-preserving mechanisms. Furthermore, while user-level DP does offer a degree of group-level privacy, its definition is fundamentally distinct from that of group privacy. Specifically, user-level DP assumes that data is collected from various users, each holding m records [23, 24, 37]. Under this paradigm, user-level DP offers theoretical privacy protections for records associated with any specific user. In contrast, group privacy safeguards information about any group of m records, thereby inherently encompassing and extending beyond the scope of user-level DP. Therefore, group privacy is more general than user-level DP, and thus algorithms satisfying user-level DP (such as those in [35, 40]) may not ensure group privacy.

7 CONCLUSION

In this paper, we present a tight general RGP bound applicable to subsampled mechanisms, based on which we derive precise and closed-form group privacy guarantees for a variety of prevalent privacy-preserving mechanisms, including subsampled Gaussian, Laplace, Skellam, and Randomized Response. For the d-dimensional Laplace mechanism within the RDP framework, our refined analysis yields a significantly tighter RDP bound, which offers a multiplicative factor saving of d and may be of independent interest. Experimental results demonstrate both the tightness and a substantial improvement of our bound over existing RGP guarantees. To the best of our knowledge, this is the first work that presents tight and closed-form RGP guarantees for subsampled mechanisms.

Regarding future work, we plan to further refine the RGP bound for subsampled Randomized Response mechanisms, and to derive closed-form RGP guarantees for other types of subsampled mechanisms such as DPIS [53], in which each record has a different probability of being included in the sample set.

ACKNOWLEDGEMENT

The authors would like to thank the anonymous reviewers for their helpful comments. This research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG3-RP-2022-029). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

Yin Yang is supported by the Qatar National Research Fund Qatar Foundation (Number NPRP11C-1229-170007). Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not reflect the views of the funding agencies.

REFERENCES

[1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *CCS*. 308–318.

[2] Naman Agarwal, Peter Kairouz, and Ziyu Liu. 2021. The skellam mechanism for differentially private federated learning. In *NeurIPS*. 5052–5064.

[3] Ritesh Ahuja, Sepanta Zeighami, Gabriel Ghinita, and Cyrus Shahabi. 2023. A Neural Approach to Spatio-Temporal Data Release with User-Level Differential Privacy. *SIGMOD* 1, 1 (2023), 1–25.

[4] Apple Differential Privacy Team. 2017. Learning with privacy at scale. http://machinelearning.apple.com/2017/12/06/learning-with-privacy-at-scale.html. Online; accessed 18 February 2022.

[5] Borja Balle, Gilles Barthe, and Marco Gaboardi. 2018. Privacy amplification by subsampling: tight analyses via couplings and divergences. In *NeurIPS*. 6280–6290.

[6] Borja Balle, Gilles Barthe, Marco Gaboardi, Justin Hsu, and Tetsuya Sato. 2020. Hypothesis testing interpretations and rényi differential privacy. In *AISTATS*. 2496–2506.

[7] Ergute Bao, Yizheng Zhu, Xiaokui Xiao, Yin Yang, Beng Chin Ooi, Benjamin Hong Meng Tan, and Khin Mi Mi Aung. 2022. Skellam mixture mechanism: a novel approach to federated learning with differential privacy. *PVLDB* 15, 11 (2022), 2348–2360.

[8] Johes Bater, Yongjoo Park, Xi He, Xiao Wang, and Jennie Rogers. 2020. Saqe: practical privacy-preserving approximate query processing for data federations. *PVLDB* 13, 12 (2020), 2691–2705.

[9] Heinz Bauer. 1958. Minimalstellen von funktionen und extremalpunkte. *Archiv der Mathematik* 9, 4 (1958), 389–393.

[10] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. 2017. Practical secure aggregation for privacy-preserving machine learning. In *CCS*. 1175–1191.

[11] Mark Bun, Marco Gaboardi, Max Hopkins, Russell Impagliazzo, Rex Lei, Toniann Pitassi, Satchit Sivakumar, and Jessica Sorrell. 2023. Stability is stable: Connections between replicability, privacy, and adaptive generalization. In *STOC*. 520–527.

[12] Clément L Canonne, Gautam Kamath, and Thomas Steinke. 2020. The discrete gaussian for differential privacy. In *NeurIPS*. 15676–15688.

[13] Kamalika Chaudhuri, Jacob Imola, and Ashwin Machanavajjhala. 2019. Capacity bounded differential privacy. In *NeurIPS*.

[14] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. 2017. Collecting telemetry data privately. In *NeurIPS*. 3574–3583.

[15] Zeyu Ding, Yuxin Wang, Danfeng Zhang, and Daniel Kifer. 2019. Free gap information from the differentially private sparse vector and noisy max mechanisms. *PVLDB* 13, 3 (2019), 293–306.

[16] Irit Dinur and Kobbi Nissim. 2003. Revealing information while preserving privacy. In *PODS*. 202–210.

[17] Wei Dong, Juanru Fang, Ke Yi, Yuchao Tao, and Ashwin Machanavajjhala. 2022. R2t: Instance-optimal truncation for differentially private query evaluation with foreign keys. In *SIGMOD*. 759–772.

[18] Cynthia Dwork. 2006. Differential privacy. In *ICALP*. 1–12.

[19] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. 2006. Our data, ourselves: Privacy via distributed noise generation. In *EUROCRYPT*. 486–503.

[20] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *TCC*. 265–284.

[21] Cynthia Dwork, Aaron Roth, et al. 2014. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* 9, 3–4 (2014), 211–407.

[22] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. 2014. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *CCS*. 1054–1067.

[23] Juanru Fang and Ke Yi. 2024. Privacy Amplification by Sampling under User-level Differential Privacy. *SIGMOD* 2, 1 (2024), 1–26.

[24] Badih Ghazi, Pritish Kamath, Ravi Kumar, Pasin Manurangsi, Raghu Meka, and Chiyuan Zhang. 2023. User-Level Differential Privacy With Few Examples Per User. In *NeurIPS*.

[25] Hongsheng Hu, Zoran Salcic, Lichao Sun, Gillian Dobbie, Philip S Yu, and Xuyun Zhang. 2022. Membership inference attacks on machine learning: A survey. *ACM Computing Surveys (CSUR)* 54, 11s (2022), 1–37.

[26] Yangfan Jiang, Xinjian Luo, Yuncheng Wu, Xiaokui Xiao, and Beng Chin Ooi. 2024. Protecting Label Distribution in Cross-Silo Federated Learning. In *IEEE S&P*. 113–113.

[27] Yangfan Jiang, Xinjian Luo, Yuncheng Wu, Xiaochen Zhu, Xiaokui Xiao, and Beng Chin Ooi. 2024. On Data Distribution Leakage in Cross-Silo Federated Learning. *IEEE TKDE* 36, 7 (2024), 3312–3328.

[28] Yangfan Jiang, Xinjian Luo, Yin Yang, and Xiaokui Xiao. 2024. Calibrating noise for group privacy in subsampled mechanisms. *arXiv preprint arXiv:2408.09943* (2024).

[29] Daniel Kifer and Ashwin Machanavajjhala. 2011. No free lunch in data privacy. In *SIGMOD*. 193–204.

[30] Daniel Kifer and Ashwin Machanavajjhala. 2012. A rigorous and customizable framework for privacy. In *PODS*. 77–88.

[31] Daniel Kifer and Ashwin Machanavajjhala. 2014. Pufferfish: A framework for mathematical privacy definitions. *TODS* 39, 1 (2014), 1–36.

[32] Ios Kotsogiannis, Yuchao Tao, Xi He, Maryam Fanaeepour, Ashwin Machanavajjhala, Michael Hay, and Gerome Miklau. 2019. Privatesql: a differentially private sql query engine. *PVLDB* 12, 11 (2019), 1371–1384.

[33] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. (2009).

[34] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.

[35] Daniel Levy, Ziteng Sun, Kareem Amin, Satyen Kale, Alex Kulesza, Mehryar Mohri, and Ananda Theertha Suresh. 2021. Learning with user-level privacy. In *NeurIPS*. 12466–12479.

[36] Junxu Liu, Jian Lou, Li Xiong, Jinfei Liu, and Xiaofeng Meng. 2021. Projected federated averaging with heterogeneous differential privacy. *PVLDB* 15, 4 (2021), 828–840.

[37] Yuhan Liu, Ananda Theertha Suresh, Felix Xinnan X Yu, Sanjiv Kumar, and Michael Riley. 2020. Learning discrete distributions: user vs item-level privacy. In *NeurIPS*. 20965–20976.

[38] Xinjian Luo, Yangfan Jiang, Fei Wei, Yuncheng Wu, Xiaokui Xiao, and Beng Chin Ooi. 2024. Exploring Privacy and Fairness Risks in Sharing Diffusion Models: An Adversarial Perspective. *IEEE TIFS* 19 (2024), 8109–8124.

[39] Miti Mazmudar, Thomas Humphries, Jiaxiang Liu, Matthew Rafuse, and Xi He. 2022. Cache Me If You Can: Accuracy-Aware Inference Engine for Differentially Private Data Exploration. *PVLDB* 16, 4 (2022), 574–586.

[40] H Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. 2018. Learning Differentially Private Recurrent Language Models. In *ICLR*.

[41] Frank D McSherry. 2009. Privacy integrated queries: an extensible platform for privacy-preserving data analysis. In *SIGMOD*. 19–30.

[42] Ilya Mironov. 2017. Rényi differential privacy. In *CSF*. 263–275.

[43] Ilya Mironov, Kunal Talwar, and Li Zhang. 2019. Rényi differential privacy of the sampled gaussian mechanism. *arXiv preprint arXiv:1908.10530* (2019).

[44] Alfréd Rényi. 1961. On measures of entropy and information. In *Berkeley symposium on mathematical statistics and probability*, Vol. 1. 547–562.

[45] Maria Rigaki and Sebastian Garcia. 2023. A Survey of Privacy Attacks in Machine Learning. *ACM Computing Surveys (CSUR)* 56 (2023), 1–34. Issue 4.

[46] Jan Schuchardt, Mihail Stoian, Arthur Kosmala, and Stephan Günnemann. 2024. Unified Mechanism-Specific Amplification by Subsampling and Group Privacy Amplification. In *NeurIPS*.

[47] Shuang Song, Yizhen Wang, and Kamalika Chaudhuri. 2017. Pufferfish privacy mechanisms for correlated data. In *SIGMOD*. 1291–1306.

[48] Florian Tramer and Dan Boneh. 2021. Differentially Private Learning Needs Better Features (or Much More Data). In *ICLR*.

[49] Salil Vadhan. 2017. The complexity of differential privacy. *Tutorials on the Foundations of Cryptography: Dedicated to Oded Goldreich* (2017), 347–450.

[50] Tim Van Erven and Peter Harremos. 2014. Rényi divergence and Kullback-Leibler divergence. *IEEE TIT* 60, 7 (2014), 3797–3820.

[51] Yu-Xiang Wang, Borja Balle, and Shiva Prasad Kasiviswanathan. 2019. Subsampled rényi differential privacy and analytical moments accountant. In *AISTATS*. 1226–1235.

[52] Stanley L Warner. 1965. Randomized response: A survey technique for eliminating evasive answer bias. *J. Amer. Statist. Assoc.* 60, 309 (1965), 63–69.

[53] Jianxin Wei, Ergute Bao, Xiaokui Xiao, and Yin Yang. 2022. Dpis: An enhanced mechanism for differentially private sgd with importance sampling. In *CCS*. 2885–2899.

[54] Yuncheng Wu, Naili Xing, Gang Chen, Tien Tuan Anh Dinh, Zhaojing Luo, Beng Chin Ooi, Xiaokui Xiao, and Meihui Zhang. 2023. Falcon: A Privacy-Preserving and Interpretable Vertical Federated Learning System. *PVLDB* 16, 10 (2023), 2471–2484.

[55] Han Xiao, Kashif Rasul, and Roland Vollgraf. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747* (2017).

[56] Hanshen Xiao, Zihang Xiang, Di Wang, and Srinivas Devadas. 2023. A Theory to Instruct Differentially-Private Learning via Clipping Bias Reduction. In *IEEE S&P*. 2170–2189.

[57] Andrew C Yao. 1982. Protocols for secure computations. In *FOCS*. 160–164.

[58] Jun Zhang, Zhenjie Zhang, Xiaokui Xiao, Yin Yang, and Marianne Winslett. 2012. Functional mechanism: regression analysis under differential privacy. In *PVLDB*. 1364–1375.

[59] Yanping Zhang, Johes Bater, Kartik Nayak, and Ashwin Machanavajjhala. 2023. Longshot: Indexing growing databases using MPC and differential privacy. *PVLDB* 16, 8 (2023), 2005–2018.

[60] Xinjing Zhou, Lidan Shou, Ke Chen, Wei Hu, and Gang Chen. 2019. DPTree: differential indexing for persistent memory. *PVLDB* 13, 4 (2019), 421–434.

[61] Yuqing Zhu and Yu-Xiang Wang. 2019. Poission subsampled rényi differential privacy. In *ICML*. 7634–7642.