

Universal Exact Compression of Differentially Private Mechanisms

Yanxiao Liu, Wei-Ning Chen, Ayfer Özgür and Cheuk Ting Li

NeurIPS 2024

Overview

Background

- In modern data science, large amounts of high-quality data are generated with personal information, often by edge devices (e.g., your messages, photos and videos).
- The data are susceptible to privacy breaches.
- Differential privacy (Warner (1965); Dwork et al. (2006)) is a powerful tool for safeguarding users' privacy by properly randomizing the local data.
- Apart from privacy, communicating (high-dimensional) local data to the central server is often a bottleneck in the system pipeline.

Kasiviswanathan, S. P., Lee, H. K., Nissim, K., Raskhodnikova, S., & Smith, A. (2011). What can we learn privately?. *SIAM Journal on Computing*, 40(3), 793-826.

Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). 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* (pp. 265-284).

Objective

We intend to answer the following fundamental question:

Practical perspective

How can we efficiently communicate privatized data?

Information-theoretic perspective

How can we exactly-compress differentially private mechanisms (which can be viewed as a noisy channel)?

Differential Privacy (DP)

Definition: Differentially Private Mechanisms)

Given a mechanism \mathcal{A} which induces distribution $P_{Z|X}$ of $Z = \mathcal{A}(X)$, we say that it satisfies (ϵ, δ) -DP if for any neighboring^a $(x, x') \in \mathcal{N}$ and $\mathcal{S} \subseteq \mathcal{Z}$, it holds that

$$\mathbf{P}(Z \in \mathcal{S} | X = x) \leq e^\epsilon \mathbf{P}(Z \in \mathcal{S} | X = x') + \delta. \quad (1)$$

^aWe say x, x' are neighboring if they differ in a single data point.

Remarks

- If a mechanism satisfies $(\epsilon, 0)$ -DP, we simply write it as ϵ -DP.

Kasiviswanathan, S. P., Lee, H. K., Nissim, K., Raskhodnikova, S., & Smith, A. (2011). What can we learn privately?. *SIAM Journal on Computing*, 40(3), 793-826.

Dwork, C., McSherry, F., Nissim, K., & Smith, A. (2006). 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* (pp. 265-284).

Related Works

Compression of DP Mechanisms

To compress ϵ -DP mechanisms:

- For $\epsilon \leq 1$, Bassily and Smith (2015) showed that a single bit can simulate any local DP randomizer with a small degradation of utility.
- Bun et al. (2019) proposed a rejection-sampling-based compression technique, which compresses an ϵ -DP mechanism into a 10ϵ -DP mechanism.
- Feldman and Talwar (2021) proposed a distributed simulation approach using rejection sampling with shared randomness.
- In Triastcyn et al. (2021); Shah et al. (2022), importance sampling (or more specifically, minimum random coding (Havasi et al. (2018))) was utilized.
- All these methods are approximate, i.e., the output distribution is distorted.

Bassily, R., & Smith, A. (2015, June). Local, private, efficient protocols for succinct histograms. In Proceedings of the forty-seventh annual ACM symposium on Theory of computing (pp. 127-135).

Bun, M., Nelson, J., & Stemmer, U. (2019). Heavy hitters and the structure of local privacy. *ACM Transactions on Algorithms (TALG)*.

Feldman, V., & Talwar, K. (2021, July). Lossless compression of efficient private local randomizers. In International Conference on Machine Learning (pp. 3208-3219). PMLR.

Shah, A., Chen, W. N., Balle, J., Kairouz, P., & Theis, L. (2022, May). Optimal compression of locally differentially private mechanisms. In International Conference on Artificial Intelligence and Statistics (pp. 7680-7723). PMLR.

Triastcyn, A., Reisser, M., & Louizos, C. (2021). Dp-rec: Private & communication-efficient federated learning. arXiv:2111.05454.

Havasi, M., Peharz, R., & Hernández-Lobato, J. M. (2018). Minimal random code learning: Getting bits back from compressed model parameters. arXiv preprint arXiv:1810.00440.

Related Works: Minimal Random Coding

In Triastcyn et al. (2021); Shah et al. (2022), importance sampling (or more specifically, minimum random coding (MRC) (Havasi et al. (2018))) was utilized.

Compressing any ϵ -DP Mechanisms by MRC (Shah et al. (2022))

- Consider an ϵ -DP mechanism $P_{Z|X}$.
- Draw $(Z_i)_{i=1,2,\dots} \stackrel{\text{iid}}{\sim} Q$. Q is known to both the encoder and the decoder^a.
- User (encoder) transmits $K \in [1 : N]$ to the server (decoder). K is drawn according to distribution $P_{\text{mrc}}(\cdot)$ such that Z_K follows $P_{Z|X}$ **approximately**.
- The distribution P_{mrc} is:
 - $\forall k \in [1 : N]$, $P_{\text{mrc}}(k) \propto w(k)$, where $w(k) := \frac{P_{Z|X}(Z_k|X)}{Q(Z_k)}$ are the importance weights.
- To communicate the index K , $\log N$ bits are required.

^aThis can be achieved via a pseudorandom number generator with a known seed.

Shah, A., Chen, W. N., Balle, J., Kairouz, P., & Theis, L. (2022, May). Optimal compression of locally differentially private mechanisms. In International Conference on Artificial Intelligence and Statistics (pp. 7680-7723). PMLR.

Triastcyn, A., Reisser, M., & Louizou, C. (2021). Dp-rec: Private & communication-efficient federated learning. arXiv:2111.05454.

Havasi, M., Peharz, R., & Hernández-Lobato, J. M. (2018). Minimal random code learning: Getting bits back from compressed model parameters. arXiv preprint arXiv:1810.00440.

Related Works: Minimal Random Coding

Algorithm 1 MRC for Compressing ϵ -LDP Mechanism $P_{Z|X}$

Input: x , ϵ -DP mechanism $P_{Z|X}$, reference distribution Q , number of candidates N

Output: compressed mechanism P_{mrc}

- 1: Draw samples $(Z_i)_{i=1,\dots,N}$ from Q .
 - 2: **for** $k = 1, 2, \dots, N$ **do**
 - 3: $w(k) \leftarrow P_{Z|X}(Z_k|x)/Q(Z_k)$
 - 4: $P_{\text{mrc}}(\cdot) \leftarrow w(\cdot)/\sum_k w(k)$
-

Theorem: If the mechanism $P_{Z|X}$ is ϵ -DP, then the compressed P_{mrc} is 2ϵ -DP.

Proof:

- For neighbors x, x' and z , since $P_{Z|X}$ is ϵ -DP, $P_{Z|X}(z|x) \leq e^\epsilon P_{Z|X}(z|x')$.
- For any $(Z_k)_{k=1,\dots,N}$ and $k \in [1 : N]$, we have

$$\begin{aligned} \frac{P_x^{\text{mrc}}(k)}{P_{x'}^{\text{mrc}}(k)} &= \frac{P_{Z|X}(Z_k|x)}{P_{Z|X}(Z_k|x')} \times \frac{\sum_i P_{Z|X}(Z_i|x')/Q(Z_i)}{\sum_i P_{Z|X}(Z_i|x)/Q(Z_i)} \\ &\leq e^\epsilon \times \frac{\sum_i e^\epsilon P_{Z|X}(Z_i|x)/Q(Z_i)}{\sum_i P_{Z|X}(Z_i|x)/Q(Z_i)} \\ &= e^{2\epsilon} \end{aligned}$$

Channel Simulation

As discussed above: all the mentioned existing methods are approximate (the output distribution is distorted). How can we exactly compress DP Mechanisms?

Channel Simulation

One-shot channel simulation (a lossy compression task) aims to find the minimal needed communication over a noiseless channel to “simulate” a channel $P_{Z|X}$.

Channel Simulation

Channel Simulation

One-shot channel simulation (a lossy compression task) aims to find the minimal needed communication over a noiseless channel to “simulate” a channel $P_{Z|X}$.

Setting

- Suppose Alice and Bob share unlimited common randomness W .
- Alice observes $X \sim P_X$ and sends a description M to Bob via a noiseless channel such that Bob can generate Z (from M and W) according to a prescribed conditional distribution $P_{Z|X}$.
- The goal is to find the minimum expected description length of M . By $X \leftrightarrow M \leftrightarrow Z$ conditional on W ,

$$\mathbb{E}[\text{Len}(M)] \geq H(M|W) \geq I(X; Z|W) = I(X; Z, W) - I(X; W) \geq I(X; Z)$$

C. H. Bennett, P. W. Shor, J. A. Smolin, and A. V. Thapliyal. Entanglement-assisted capacity of a quantum channel and the reverse Shannon theorem. *IEEE Trans. Inf. Theory*, vol. 48, no. 10, pp. 2637–2655, Oct. 2002.

Harsha, P., Jain, R., mathcallester, D., & Radhakrishnan, J. (2007, June). The communication complexity of correlation. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)* (pp. 10-23). IEEE.

P. Cuff. Distributed channel synthesis. *IEEE Trans. Inf. Theory*, vol. 59, no. 11, pp. 7071–7096, Nov. 2013.

C. H. Bennett, I. Devetak, A. W. Harrow, P. W. Shor, and A. Winter. The quantum reverse Shannon theorem and resource tradeoffs for simulating quantum channels. *IEEE Trans. Inf. Theory*, vol. 60, no. 3, pp. 2926–2959, May 2014.

Channel Simulation

Channel Simulation

- By Harsha et al. (2007) and Li and El Gamal (2018), $P_{Z|X}$ can be simulated using $I(X; Z) + O(\log(I(X; Z)))$ bits.
- In Harsha et al. (2007), algorithms based on rejection sampling are proposed.
- Dithered quantization (Ziv (1985)) has been used to simulate an additive noise channel in Agustsson and Theis (2020) for neural compression.
- More applications of channel simulation tools:
 - Neural network compression by Havasi et al. (2018)
 - Image compression via variational autoencoders by Flamich et al. (2020)
 - Diffusion models with perfect realism by Theis et al. (2022)
 - Differentially private federated learning by Shah et al. (2022)

Poisson Private Representation (PPR)

Poisson Private Representation: Overview(PPR)

- In this paper, we propose a novel algorithm, called **Poisson private representation** (PPR), that is designed to compress and simulate any local randomizer while ensuring local differential privacy.
- The advantages of our PPR are as follows:
 - ① **Universality**: Unlike dithered-quantization-based approaches which can only simulate additive noise mechanisms, PPR can simulate any local or central DP mechanism with discrete or continuous input and output.
 - ② **Exactness**: PPR enables exact simulation, ensuring that the reproduced distribution perfectly matches the original one, and hence the compressed sample maintains the same statistical properties.
 - ③ **Communication efficiency**: PPR compresses the output of any DP mechanism to a size close to the theoretical lower bound $I(X; Z)$.
- PPR is the first universal exact compression method for DP mechanisms with an almost-optimal compression size.
 - The methods by Bassily and Smith (2015); Bun et al. (2019); Feldman and Talwar (2021); Shah et al. (2022) are not exact.
 - The methods by Harsha et al. (2007) and Li and El Gamal (2018) do not guarantee privacy.

Poisson Functional Representation (PFR)

To design our Poisson Private Representation, we first review the Poisson Functional Representation by Li and El Gamal (2018):

Definition

Let $(T_i)_i$ be a Poisson process with rate 1 (i.e., $T_1, T_2 - T_1, T_3 - T_2, \dots \stackrel{\text{iid}}{\sim} \text{Exp}(1)$), independent of $Z_i \stackrel{\text{iid}}{\sim} Q$ for $i = 1, 2, \dots$. Then $(Z_i, T_i)_i$ is a Poisson process with intensity measure $Q \times \lambda_{[0, \infty)}$, where $\lambda_{[0, \infty)}$ is the Lebesgue measure over $[0, \infty)$. Fix any distribution P over \mathcal{Z} that is absolutely continuous with respect to Q . Let

$$\tilde{T}_i := T_i \cdot \left(\frac{dP}{dQ}(Z_i) \right)^{-1}. \quad (2)$$

Then $(Z_i, \tilde{T}_i)_i$ is a Poisson process with intensity measure $P \times \lambda_{[0, \infty)}$. The **Poisson functional representation** selects $Z = Z_K$ with the smallest associated \tilde{T}_K , i.e., let $K := \operatorname{argmin}_i \tilde{T}_i$ and $Z := Z_K$.

Li, C. T., & El Gamal, A. (2018). Strong functional representation lemma and applications to coding theorems. IEEE Transactions on Information Theory.

Poisson Functional Representation (PFR)

Poisson Functional Representation

- The PFR selects a sample following the target distribution P using another distribution Q .
- It draws a random sequence $(Z_i)_i$ from Q and a sequence of times (T_i) according to a Poisson process.
- If we select Z_i with the smallest T_i , the selected sample follows Q .
- To obtain a sample from P instead, we multiply the time by the factor $\left(\frac{dP}{dQ}(Z_i)\right)^{-1}$ so the Z_i with the smallest \tilde{T}_i .
- However, if one uses PFR to simulate a DP mechanism $P_{Z|X}$, with inputs $X, (T_i)_i$ and output K , the issue is that K is a function of X and the shared randomness $(Z_i, T_i)_i$.
- Hence a change of X may affect K in a deterministic manner, and hence the privacy of X cannot be protected well.

Li, C. T., & El Gamal, A. (2018). Strong functional representation lemma and applications to coding theorems. IEEE Transactions on Information Theory.

Poisson Private Representation (PPR)

Poisson Private Representation: Construction

Input: x , (ϵ, δ) -DP mechanism $P_{Z|X}$, reference distribution Q , parameter $\alpha > 1$.

- 1 Generate shared randomness between user and server $(Z_i)_{i=1,2,\dots} \stackrel{\text{iid}}{\sim} Q$.
- 2 The user knows $(Z_i)_i$, x , $P_{Z|X}$ and performs:
 - 1 Generate the Poisson process $(T_i)_i$ with rate 1.
 - 2 Compute $\tilde{T}_i := T_i \cdot \left(\frac{dP}{dQ}(Z_i)\right)^{-1}$ where $P := P_{Z|X}(\cdot|x)$. Take $\tilde{T}_i = \infty$ if $\frac{dP}{dQ}(Z_i) = 0$.
 - 3 Generate $K \in \mathbb{Z}_+$ with

$$\mathbf{P}(K = k) = \frac{\tilde{T}_k^{-\alpha}}{\sum_{i=1}^{\infty} \tilde{T}_i^{-\alpha}}. \quad (3)$$

- 4 Compress and send K (e.g., by Elias delta code).
- 3 The server, which observes $(Z_i)_i$ and K , outputs $Z = Z_K$.

Remarks

- While the algorithm requires infinite samples, it can be reparametrized to terminate in finite steps.
- When $\alpha = \infty$, PPR reduces to PFR.

Poisson Private Representation (PPR): Theoretic Guarantees

Proposition: Exactness

The output Z of PPR follows $P_{Z|X}$ exactly.

Remarks

- Due to the **exactness** of PPR, it guarantees unbiasedness for tasks such as distributed mean estimation:
 - 1 If we only want a stand-alone privacy mechanism, we can just focus on the privacy and utility.
 - 2 However, if the output of the mechanism is used for downstream tasks (e.g., after receiving information from clients, the server sends aggregated mean to data analysts), having an exact distribution allows more precise (central) privacy and utility guarantees.
- PPR is **universal**: only the encoder needs to know the simulated mechanism $P_{Z|X}$.

Poisson Private Representation (PPR): Theoretic Guarantees

Theorems: Privacy Guarantee

- 1 **Theorem 4.5** (ϵ -DP of PPR): If the mechanism $P_{Z|X}$ is ϵ -DP, then PPR $P_{(Z_i)_{i,K}|X}$ with parameter $\alpha > 1$ is $2\alpha\epsilon$ -DP.
- 2 **Theorem 4.6** ((ϵ, δ) -DP of PPR): If $P_{Z|X}$ is (ϵ, δ) -DP, then PPR $P_{(Z_i)_{i,K}|X}$ with parameter $\alpha > 1$ is $(2\alpha\epsilon, 2\delta)$ -DP.

Theorems: Privacy Guarantee

Theorem 4.8 (Tighter (ϵ, δ) -DP of PPR): If $P_{Z|X}$ is (ϵ, δ) -DP, then PPR $P_{(Z_i)_{i,K}|X}$ with parameter $\alpha > 1$ is $(\alpha\epsilon + \tilde{\epsilon}, 2(\delta + \tilde{\delta}))$ -DP, for every $\tilde{\epsilon} \in (0, 1]$ and $\tilde{\delta} \in (0, 1/3]$ that satisfy $\alpha \leq e^{-4.2\tilde{\delta}\tilde{\epsilon}^2}/(-\ln \tilde{\delta}) + 1$.

Privacy Guarantee on Compressing ϵ -DP Mechanisms

- If the mechanism $P_{Z|X}$ is ϵ -DP, then PPR $P_{(Z_i)_i, K|X}$ is $2\alpha\epsilon$ -DP for $\alpha > 1$.
- Proof: For $j = 1, 2$,

- for neighbor x_1, x_2 , suppose $P_j := P_{Z|X}(\cdot|x_j)$ and $\tilde{T}_{j,i} := T_i / (\frac{dP_j}{dQ}(Z_i))$.
- Since $P_{Z|X}$ is ϵ -DP, we know

$$e^{-\epsilon} \frac{dP_2}{dQ}(z) \leq \frac{dP_1}{dQ}(z) \leq e^{\epsilon} \frac{dP_2}{dQ}(z).$$

and hence $e^{-\epsilon} \tilde{T}_{2,i} \leq \tilde{T}_{1,i} \leq e^{\epsilon} \tilde{T}_{2,i}$.

- Suppose K_j is PPR's output applied on P_j . For $k \in \mathbb{Z}_+$, almost surely,

$$\mathbf{P}(K_1 = k | (Z_i, T_i)_i) = \frac{\tilde{T}_{1,k}^{-\alpha}}{\sum_{i=1}^{\infty} \tilde{T}_{1,i}^{-\alpha}} \leq \frac{e^{\alpha\epsilon} \tilde{T}_{2,k}^{-\alpha}}{\sum_{i=1}^{\infty} e^{-\alpha\epsilon} \tilde{T}_{1,i}^{-\alpha}} = e^{2\alpha\epsilon} \mathbf{P}(K_2 = k | (Z_i, T_i)_i).$$

- For any measurable $S \subseteq \mathcal{Z}^{\infty} \times \mathbb{Z}_{\geq 0}$,

$$\begin{aligned} \mathbf{P}(((Z_i)_i, K_1) \in S) &= \mathbf{E}[\mathbf{P}(((Z_i)_i, K_1) \in S | (Z_i, T_i)_i)] \\ &= \mathbf{E} \left[\sum_{k: ((Z_i)_i, k) \in S} \mathbf{P}(K_1 = k | (Z_i, T_i)_i) \right] \\ &\leq e^{2\alpha\epsilon} \mathbf{E} \left[\sum_{k: ((Z_i)_i, k) \in S} \mathbf{P}(K_2 = k | (Z_i, T_i)_i) \right] \\ &= e^{2\alpha\epsilon} \mathbf{P}(((Z_i)_i, K_2) \in S) \end{aligned}$$

Poisson Private Representation (PPR): Theory

Theorem: Communication efficiency

Theorem 4.3 (Compression size of PPR): For PPR with parameter $\alpha > 1$, message K satisfies

$$\mathbf{E}[\log_2 K] \leq D_{\text{KL}}(P(\cdot|x) \| Q(\cdot)) + \frac{\log_2(3.56)}{\min\{\frac{\alpha-1}{2}, 1\}}.$$

As a result, when the input $X \sim P_X$ is random, we have

$$\mathbf{E}[\log_2 K] \leq I(X; Z) + \frac{\log_2(3.56)}{\min\{\frac{\alpha-1}{2}, 1\}}.$$

Hence, K can be encoded into $I(X; Z) + \log_2(I(X; Z) + 1) + O(1)$ bits, close to the theoretical lower bound $I(X; Z)$.

Running Time

Running Time

- Note since $\mathbf{E}[\log K] \approx I(X; Z)$, K (and hence the running time) is at least exponential in $I(X; Z)$.
- However, an exponential complexity is also needed in sampling methods without privacy guarantee, e.g., Havasi et al. (2018) and Maddison (2016).
- By Agustsson and Theis (2020), no polynomial time general sampling-based method exists (even without privacy constraint), if $RP \neq NP$.
- Nevertheless, this is not an obstacle: $I(X; Z)$ for a good local DP mechanism must be small, or the leakage of X in Z will be too large^a.
- Another way to ensure a polynomial running time is to divide the data into small chunks and apply the mechanism to each chunk separately.

^aBy Cuff and Yu (2016), for an ϵ -local DP mechanism, $I(X; Z) \leq \min\{\epsilon, \epsilon^2\}$

Chris J Maddison. A Poisson process model for Monte Carlo. Perturbation, Optimization, and Statistics, pages 193–232, 2016.

Marton Havasi, Robert Peharz, and José Miguel Hernández-Lobato. Minimal random code learning: Getting bits back from compressed model parameters. In 7th International Conference on Learning Representations, ICLR 2019, 2019.

Eirikur Agustsson and Lucas Theis. Universally quantized neural compression. Advances in neural information processing systems, 33:12367–12376, 2020.

Paul Cuff and Lanqing Yu. Differential privacy as a mutual information constraint. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pages 43–54, 2016.

Application: Distributed Mean Estimation (DME)

Distributed Mean Estimation

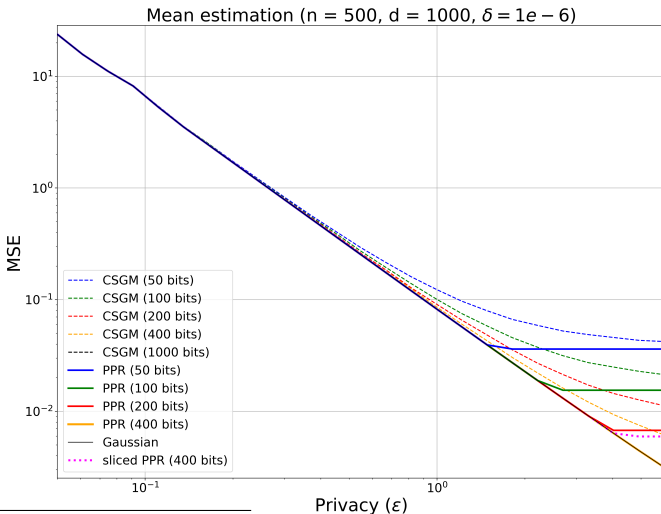
- Private DME is the core sub-routine in various private and federated optimization algorithms.
- Consider n users, each with data $X_i \in \mathbb{R}^d$.
- They use **Gaussian mechanism** and send $Z_i \sim \mathcal{N}\left(X_i, \frac{\sigma^2}{n} \mathbb{I}_d\right)$ to server, where $\sigma \geq \frac{C\sqrt{2\ln(1.25/\delta)}}{\epsilon}$. Server estimates the mean as $\hat{\mu}(Z^n) = \frac{1}{n} \sum_i Z_i$.
- Using PPR to compress the Gaussian mechanism:
 - $\hat{\mu}(Z^n) = \frac{1}{n} \sum_i Z_i$ is unbiased, has (ϵ, δ) -central DP.
 - PPR satisfies $(2\alpha\sqrt{n}\epsilon, 2\delta)$ -local DP for $\epsilon < \frac{1}{\sqrt{n}}$.
 - The average per-client communication cost is at most $\ell + \log_2(\ell + 1) + 2$ bits, where

$$\ell := \frac{d}{2} \log_2 \left(\frac{C^2 n}{d\sigma^2} + 1 \right) + \eta_\alpha \leq \frac{d}{2} \log_2 \left(\frac{n\epsilon^2}{2d \ln(1.25/\delta)} + 1 \right) + \eta_\alpha,$$

$$\text{and } \eta_\alpha := (\log_2(3.56)) / \min\{(\alpha - 1)/2, 1\}.$$
^a


^aThis is at least as good as $O\left(\frac{n\epsilon^2}{\log(1/\delta)+1}\right)$ in Suresh et al. (2017) and Chen et al. (2024), and is better when $n \gg d$.

Application: Distributed Mean Estimation



Chen, W. N., Song, D., Ozgur, A., & Kairouz, P. (2024). Privacy amplification via compression: Achieving the optimal privacy-accuracy-communication trade-off in distributed mean estimation. *Advances in Neural Information Processing Systems*, 36.

Application: Distributed Mean Estimation



- PPR consistently achieves a smaller MSE compared to CSGM.
- The MSE reductions are significant:
 - For $\epsilon = 1$ and compressing to 50 bits, we give a 33.61% MSE reduction;
 - For $\epsilon = 0.5$ and compressing to 25 bits, we give a 22.33% MSE reduction;
 - All considered mechanisms are asymptotically close to optimal.
- We emphasize that PPR provides both central and local DP guarantees.
- We ensure $O(d)$ running time by breaking the vector into small chunks (slided PPR with 50-bits chunks).

Application: Metric Privacy and Laplace Mechanism

Metric Privacy

For a mechanism \mathcal{A} with $P_{Z|X}$ and a metric $d_{\mathcal{X}}$ over \mathcal{X} , it satisfies $\epsilon \cdot d_{\mathcal{X}}$ -privacy (Andrés et al. (2013)) if $\forall x, x' \in \mathcal{X}, \mathcal{S} \subseteq \mathcal{Z}$, we have

$$\mathbf{P}(Z \in \mathcal{S} | X = x) \leq e^{\epsilon \cdot d_{\mathcal{X}}(x, x')} \mathbf{P}(Z \in \mathcal{S} | X = x').$$

Laplace Mechanism

- Users to send privatized data Z to an untrusted server, so that the server can know X approximately but not exactly.
- Laplace mechanism (conditional density function $f_{Z|X} \propto e^{-\epsilon d_{\mathcal{X}}(x, z)}$ and $d_{\mathcal{X}}(x, z) = \|x - z\|_2$) achieves $\epsilon \cdot d_{\mathcal{X}}$ -privacy.
- It has been used in geo-indistinguishability to privatize the users' locations.

Andrés, M. E., Bordenabe, N. E., Chatzikokolakis, K., & Palamidessi, C. (2013, November). Geo-indistinguishability: Differential privacy for location-based systems. In Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security (pp. 901-914).

Application: Metric Privacy and Laplace Mechanism

In Andrés et al. (2013), a discrete Laplace mechanism has been used, where each coordinate of a real vector is quantized to a finite number of levels, introducing additional distortion.

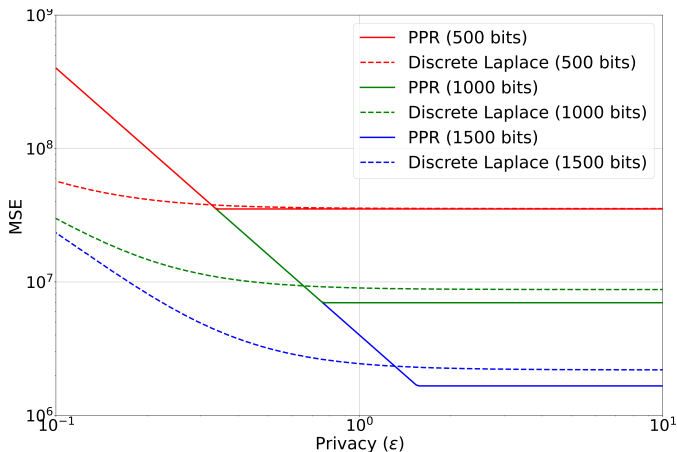
PPR-compressed Laplace mechanism

Consider PPR applied to the Laplace mechanism $P_{Z|X}$ where $X \in \{x \in \mathbb{R}^d \mid \|x\|_2 \leq C\}$ and a proposal distribution $Q = \mathcal{N}(0, (\frac{C^2}{d} + \frac{d+1}{\epsilon^2})\mathbb{I}_d)$. The output of PPR achieves an MSE $\frac{d(d+1)}{\epsilon^2}$, a $2\alpha\epsilon \cdot d_{\mathcal{X}}$ -privacy and compression size at most $\ell + \log_2(\ell + 1) + 2$ bits, where $\ell :=$

$$\frac{d}{2} \log_2 \left(\frac{2}{e} \left(\frac{C^2 \epsilon^2}{d} + d + 1 \right) \right) - \log_2 \frac{\Gamma(d+1)}{\Gamma(\frac{d}{2} + 1)} + \frac{\log_2(3.56)}{\min\{\frac{\alpha-1}{2}, 1\}}.$$

We compare with the discrete Laplace mechanism by Andrés et al. (2013).

Application: Metric Privacy and Laplace Mechanism



Andrés, M. E., Bordenabe, N. E., Chatzikokolakis, K., & Palamidessi, C. (2013, November). Geo-indistinguishability: Differential privacy for location-based systems. In Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security (pp. 901-914).

Summary

Summary

- We proposed a novel scheme for compressing DP mechanisms, called **Poisson private representation (PPR)**.
- Unlike previous schemes which are either constrained on special classes of DP mechanisms or introducing additional distortions on the output, our scheme can compress and **exactly** simulate **arbitrary** mechanisms while providing privacy guarantees.
- PPR provides a compression size that is close to the theoretic lower bound.

Future Works

- Reduce the running time of PPR under certain restrictions. For example, for unimodal $P_{Z|X}$, techniques utilized by Flamich et al. (2022); Flamich (2024) could be useful.

Flamich, G. (2024). Greedy Poisson rejection sampling. Advances in Neural Information Processing Systems, 36.

Flamich, G., Markou, S., & Hernández-Lobato, J. M. (2022, June). Fast relative entropy coding with a* coding. In International Conference on Machine Learning (pp. 6548-6577). PMLR.

Acknowledgement

Yanxiao Liu was partially supported by the CUHK PhD International Mobility for Partnerships and Collaborations Award 2023-24 and CUHK Postgraduate Scholarship. Wei-Ning Chen and Ayfer Özgür were supported by the NSF grant CIF-2213223. Cheuk Ting Li was partially supported by two grants from the Research Grants Council of the Hong Kong Special Administrative Region, China [Project No.s: CUHK 24205621 (ECS), CUHK 14209823 (GRF)].

References

- Agustsson, E. and Theis, L. (2020). Universally quantized neural compression. *Advances in neural information processing systems*, 33:12367–12376.
- Andrés, M. E., Bordenabe, N. E., Chatzikokolakis, K., and Palamidessi, C. (2013). Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 901–914.
- Bassily, R. and Smith, A. (2015). Local, private, efficient protocols for succinct histograms. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 127–135.
- Bun, M., Nelson, J., and Stemmer, U. (2019). Heavy hitters and the structure of local privacy. *ACM Transactions on Algorithms (TALG)*, 15(4):1–40.
- Chen, W.-N., Song, D., Ouzgur, A., and Kairouz, P. (2024). Privacy amplification via compression: Achieving the optimal privacy-accuracy-communication trade-off in distributed mean estimation. *Advances in Neural Information Processing Systems*, 36.
- Cuff, P. and Yu, L. (2016). Differential privacy as a mutual information constraint. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 43–54.
- Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006). Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer.
- Feldman, V. and Talwar, K. (2021). Lossless compression of efficient private local randomizers. In *International Conference on Machine Learning*, pages 3208–3219. PMLR.
- Flamich, G. (2024). Greedy poisson rejection sampling. *Advances in Neural Information Processing Systems*, 36.
- Flamich, G., Havasi, M., and Hernández-Lobato, J. M. (2020). Compressing images by encoding their latent representations with relative entropy coding. *Advances in Neural Information Processing Systems*, 33:16131–16141.
- Flamich, G., Markou, S., and Hernández-Lobato, J. M. (2022). Fast relative entropy coding with a* coding. In *International Conference on Machine Learning*, pages 6548–6577. PMLR.
- Harsha, P., Jain, R., McAllester, D., and Radhakrishnan, J. (2007). The communication complexity of correlation. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)*, pages 10–23. IEEE.
- Havasi, M., Peharz, R., and Hernández-Lobato, J. M. (2018). Minimal random code learning: Getting bits back from compressed model parameters. *arXiv preprint arXiv:1810.00440*.
- Li, C. T. and El Gamal, A. (2018). Strong functional representation lemma and applications to coding theorems. *IEEE Transactions on Information Theory*, 64(11):6967–6978.
- Maddison, C. J. (2016). A poisson process model for monte carlo. *Perturbation, Optimization, and Statistics*, pages 193–232.
- Shah, A., Chen, W.-N., Balle, J., Kairouz, P., and Theis, L. (2022). Optimal compression of locally differentially private mechanisms. In *International Conference on Artificial Intelligence and Statistics*, pages 7680–7723. PMLR.
- Suresh, A. T., Felix, X. Y., Kumar, S., and McMahan, H. B. (2017). Distributed mean estimation with limited communication. In *International conference on machine learning*, pages 3329–3337. PMLR.
- Theis, L., Salimans, T., Hoffman, M. D., and Mentzer, F. (2022). Lossy compression with gaussian diffusion. *arXiv preprint arXiv:2206.08889*.
- Triastcyn, A., Reisser, M., and Louizos, C. (2021). Dp-rec: Private & communication-efficient federated learning. *arXiv preprint arXiv:2111.05454*.
- Warner, S. L. (1965). Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American statistical association*, 60(309):63–69.