Instance-Wise Laplace Mechanism via Deep Reinforcement Learning (Student Abstract)

Sehyun Ryu, Hosung Joo, Jonggyu Jang, Hyun Jong Yang

Advanced Information Systems Lab, Dept. of Electrical Engineering, POSTECH
77, Cheongam-ro, Nam-gu, Pohang-si, Gyeongsangbuk-do
37673, Republic of Korea
{ sh.ryu, hosung.joo, jgjang, hyunyang }@postech.ac.kr

Abstract

Recent research has shown a growing interest in per-instance differential privacy (pDP), highlighting the fact that each data instance within a dataset may incur distinct levels of privacy loss. However, conventional additive noise mechanisms apply identical noise to all query outputs, thereby deteriorating data statistics. In this study, we propose an instance-wise Laplace mechanism, which adds non-identical Laplace noises to the query output for each data instance. A challenge arises from the complex interaction of additive noise, where the noise introduced to individual instances impacts the pDP of other instances, adding complexity and resilience to straightforward solutions. To tackle this problem, we introduce an instance-wise Laplace mechanism algorithm via deep reinforcement learning and validate its ability to better preserve data statistics on a real dataset, compared to the original Laplace mechanism.

Introduction

The concept of differential privacy (DP) is first introduced by Dwork (2006) for safeguarding the privacy of individual data points with the level of ϵ. The Laplace mechanism is a representative additive noise mechanism for guaranteeing ϵ-DP, which can be simply applied by adding identical Laplace noises to query outputs. However, a new definition of per-instance DP (pDP) pointed out the need for non-identical additive noise since each query output naturally has a different level of privacy (Wang et al. 2019). However, there have been no studies on non-identical noise optimization for pDP manner due to correlated privacy loss and output distributions. In this study, our focus is to propose the instance-wise Laplace mechanism via deep reinforcement learning (DRL) (Mnih et al. 2015).

Contribution Our finding is the first instance-wise Laplace mechanism. We design the optimization process as a Markov decision process (MDP). To solve the MDP problem, we propose a DQN-based algorithm aiming to guarantee ϵ-pDP and to preserve the data statistics. The numerical results ensure that the proposed method better preserve data statistics via carefully designing appropriate noise for each data instance.

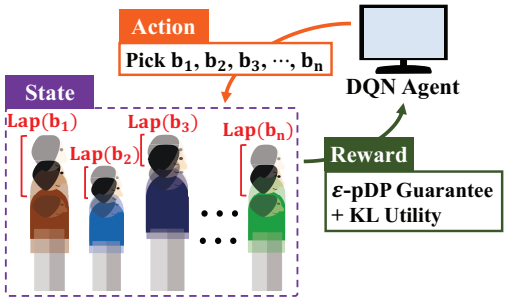

Figure 1: A deep reinforcement learning approach to find an optimal variance set for instance-wise Laplace mechanism, to better preserve data statistics while guaranteeing ϵ-pDP.

Backgrounds

Per-instance-DP For a given dataset $\mathcal{Z}$, a randomized mechanism $\mathcal{M}$ satisfies ϵ-pDP if the following inequality holds for a data point z and all $S \subseteq \text{Range}(\mathcal{M})$:

$$\Pr[\mathcal{K}(\mathcal{Z}) \in S] \leq \exp(\epsilon) \times \Pr[\mathcal{K}(\mathcal{Z} \cup \{z\}) \in S]. \quad (1)$$

Laplace mechanism Let us define the l_1 sensitivity of a function $f : \mathbb{N}^{|\mathcal{X}|} \rightarrow \mathbb{R}^k$ as $\Delta f = \max_{z_1, z_2 \in \mathcal{Z}} \|f(z_1) - f(z_2)\|_1$. Then, the Laplace mechanism for satisfying ϵ-pDP is denoted as:

$$\mathcal{M}_{\text{LM}}(x, f(\cdot), \epsilon) = f(x) + (Y_1, Y_2, ..., Y_k), \quad (2)$$

where Y_i are i.i.d random variables drawn from $\text{Lap}(\Delta f/\epsilon)$.

Random sampling query We set our target query as a random sampling query. Given numeric dataset $\mathcal{Z}$, the output of a random sampling query q is a random variable following the probability distribution of a data set:

$$q(\mathcal{Z}) \sim \Pr(\mathcal{Z}). \quad (3)$$

Methodology

We introduce the concept of applying Laplace noise with varying variances to each query output of data instances — namely, instance-wise Laplace mechanism (ILM). This approach offers the potential to better preserve data statistics compared to the conventional Laplace mechanism while maintaining the same ϵ-pDP for all data.

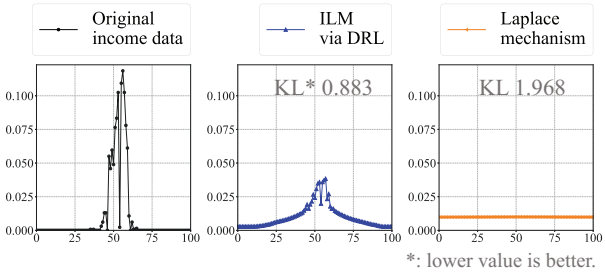

Figure 2: Comparison of query output probability distributions and KL-divergences between original distribution, for the height data with each algorithm for $\epsilon = 0.011$.

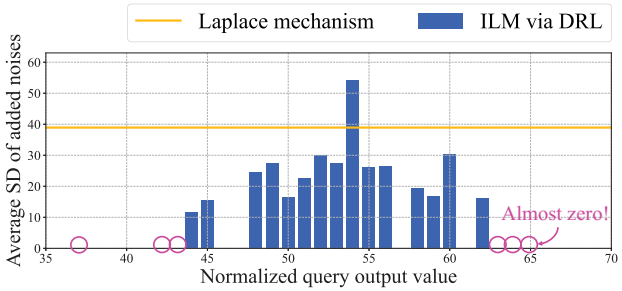

Figure 3: Comparison of chosen noise variance distributions for the height data with each algorithm for $\epsilon = 0.011$.

Definition 1 (Instance-wise Laplace mechanism) *Given any function* $f : \mathbb{N}^{|\mathcal{D}|} \rightarrow \mathbb{R}^k$ *for a fixed data set* $\mathcal{D} = \bigcup_{i=1}^{|\mathcal{D}|} d_i$*, the instance-wise Laplace mechanism is defined as:*

$$\mathcal{K}_I(f(d_i), \epsilon) = f(d_i) + (Y_{i1}, Y_{i2}, ..., Y_{ik}), \tag{4}$$

where Y_{ij} *are random variables drawn form* $\text{Lap}(b_{ij})$.

Implementing ILM in practice poses a challenge due to the strong correlation among the pDP values associated with each datum. Thus we propose to exploit deep reinforcement learning (DRL), which could find a pattern in a complex situation, to find an optimal variance set for ILM, as seen in Fig. 1. We define one episode as specifying the noise distribution of all data in the data set once.

State The state is formed by concatenating two vectors corresponding to the positional-encoded (Vaswani et al. 2017) current query output value and the PMF of total query outputs.

Action The action involves determining a noise variance value for a current data instance and adding this to its query output. In our simulation, we utilized the five options as a variance-related b value, $\frac{3\times\Delta f}{\epsilon}$, $\frac{2\times\Delta f}{\epsilon}$, $\frac{\Delta f}{2}$, $\frac{0.01\times\Delta f}{2}$, and $\frac{0.001\times\Delta f}{2}$, where the $b = \frac{\Delta f}{\epsilon}$ for the Laplace mechanism's noise.

Reward The reward of the model is composed of the sum of two values, $R = R_{\text{E}} + R_{\text{U}}$, to minimize KL-divergence while guaranteeing ϵ-pDP. R_{E} plays a role in determining whether the impact of adding noise satisfies the ϵ-pDP constraint. If the current noise satisfies ϵ-pDP, a reward 1 is given; otherwise, a penalty -1 is imposed. The R_{U} is given by $1 - \frac{D_{\text{KL}}(q(\mathcal{D})||\mathcal{A}_{\text{DRL}}(q(\mathcal{D})))}{D_{\text{KL}}(q(\mathcal{D})||\mathcal{A}_{\text{LAP}}(q(\mathcal{D})))}$, if $D_{\text{KL}}(q(\mathcal{D})||\mathcal{A}_{\text{DRL}}(q(\mathcal{D})))$ (the KL-divergence of the DRL-modified distribution) is less than $D_{\text{KL}}(q(\mathcal{D})||\mathcal{A}_{\text{LAP}}(q(\mathcal{D})))$ (the KL-divergence of the Laplace mechanism-modified distribution). Otherwise, R_{U} is given as -1.

Experiments and Results

We have evaluated our method with a personal height dataset, publicly available at Kaggle competition[1]. In our experiment, we carefully configure the smallest possible value of ϵ which is guaranteed ϵ-pDP with our algorithm.

Figure 2 compares the distributions of the original dataset and randomized dataset. As the figure shows, the proposed method much better preserves the shape of the distribution compared to the conventional mechanism, even though both methods achieve the same level of ϵ-pDP. In the case of a very small ϵ, the Laplace mechanism nearly erases the original characteristics of the dataset. To quantitatively measure the data statistics, we compare the data statistics in KL divergence utility. The proposed method achieves lower KL-divergence (0.883) compared to the conventional Laplace mechanism (1.968), which means our method quantitatively better preserves statistical utility.

Figure 3 shows the standard deviation of added noises for each data instance. As shown in the figure, the proposed method exploits noises of varying and generally lower variances distinct from the conventional Laplace mechanism. More importantly, the proposed method applies relatively low-variance noises to frequently occurring data instances, which is a natural approach because the more frequently data instances occur, the lower the privacy risks are.

Discussion and Future Directions

Our finding is the first work on the instance-wise Laplace mechanism; however, the proposed method lacks generality in various types of noise distributions such as Gaussian. In future works, we will target:

- A generalized mechanism for instance-wise pDP.
- A Low-complex pDP mechanism.

References

Dwork, C. 2006. Differential Privacy. In *Automata, Languages and Programming*, 1–12. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-540-35908-1.

Mnih; et al. 2015. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533.

Vaswani; et al. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

Wang; et al. 2019. Per-instance Differential Privacy. *Journal of Privacy and Confidentiality*, 9(1).

[1]https://www.kaggle.com/datasets/justinas/nba-players-data