Differentially Private Machine Learning

Kanaparthy S V Samhita
Machine Learning Lab,
International Institute of Information Technology,
s.v.samhita@research.iiit.ac.in

Sujit P Gujar
Machine Learning Lab,
International Institute of Information Technology,
sujit.gujar@iiit.ac.in

Abstract

This report covers the basics of Differential Privacy (DP) and Multi-Armed Bandits (MAB). It also introduces Local Differential Privacy (LDP) that is widely explored by researchers these days. Later it includes the summaries of a few recent papers on Differential Privacy and Multi-Armed Bandits problems that provide privacy guarantees using Local Differential Private algorithms.

Contents

1 Introduction **1**

2 Differential Privacy **2**
2.1 Local Differential Privacy (LDP) 2

3 Multi-Armed Bandits (MAB) **3**
3.1 Privacy in MAB Setting 3

4 RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response **4**
4.1 Problem Addressed 4
4.2 Previous Work 4
4.3 Novelty 4
4.4 RAPPOR 4
4.4.1 Randomization Algorithm 4
4.4.2 Utility Decoding 5

5 Differentially Private and Fair Deep Learning: A Lagrangian Dual Approach **5**
5.1 Problem Addressed 5
5.2 Previous Work 5
5.3 Problem Setting 5
5.4 Approach 5
5.5 Fairness 5
5.6 Privacy 6
5.7 Conclusion 6

6 Multi-Armed Bandits with Local Differential Privacy **6**
6.1 Problem Addressed 6
6.2 Model Overview 6
6.2.1 Bandit Model 6
6.2.2 Regret Minimisation 6
6.3 Previous Work 6
6.4 Lower Bound 7
6.5 Algorithms and Upper Bounds 7

7 Local differential privacy for social network publishing **7**
7.1 Problem Addressed 7
7.2 Previous Work 7
7.3 Local Differential Privacy Model 8
7.3.1 Algorithm 8

8 Principal Component Analysis in the Local Differential Privacy Model **8**
8.1 Problem Addressed 8
8.2 Approach 9
8.3 Previous Work 9
8.4 Minimax Risk 9
8.4.1 Classical Minimax Risk 9
8.4.2 Private Minimax Risk 9
8.5 Results 9

9 Overview **10**

List of Papers Summarized **10**

References **11**

1

Introduction

With the advent and increased use of the internet, social networks has become an essential part of people's daily routine. Social network is not only used to connect with others, but it has become an effective platform for businesses/marketing to reach their target audience. The emergence of big data advanced social network marketing to a new level. This made an enormous amount of data available which marketers are able to utilize to get insights for framing efficient social network marketing strategies. All the user activities like status updates, photos, and videos posted on various social networks contain useful information about their demographics, likes, dislikes. Businesses utilize this information in numerous ways, managing and analyzing it to get a competitive edge. This enables personalization allowing brands to approach their customers in a more personalized way based on their choices. It gives in-depth insights and a holistic understanding of the audience, which aids businesses in creating tailored communication for them to enhance their trust. This allows brands to easily display only those advertisements which interest consumers, turning ads into a non-obtrusive experience for the users. These advertisements will be targeted based on users' social network posts, what they watch and share, etc. With personalized advertisements, it will be possible for marketers to strengthen their relationships with social network users and convert them into customers. The more information you get about consumers, the better you will be able to target them through your social network marketing. However, this throws light on the privacy of the users on these social networks. There may be concerns over privacy as it results in the collection of personal data of individuals.

In recent years, users have become increasingly concerned about protecting their privacy online information and activities, which may include their personal profiles, browsing histories, and activities on the Internet. They may not want to share this information with other parties. However as mentioned in the previous scenario, many real-world systems like recommender systems, advertisement allocators, online shopping websites, medical experiments, and search engines need such data to learn critical matters and provide better services. To handle this dilemma, there is a compelling need to develop algorithms that can optimally trade-off system performance and the privacy level provided to the users.

A widely accepted and applied metric to measure the privacy level is the Differential Privacy (DP) [14], which, in theory, guarantees that it is difficult for any party or eavesdropper to determine whether or not an individual is listed in a private database. DP algorithms have been studied in many areas [24, 29], explained formally in following sections. [1]

2

Differential Privacy

Differential Privacy makes it possible for social networks to collect and share aggregate information about user habits, while maintaining the privacy of individual users. The formal definition of Differential Privacy is as follows,

Definition 2.1 ((ϵ, δ)-Differential Privacy (DP))**.** *For $\epsilon > 0$, a randomized mapping $M : \mathcal{D} \to R^l$ is said to be (ϵ, δ)-DP on $\mathcal{D} \subset \mathbb{R}^k$ if for any neighbouring x, x' in $\mathcal{D}$ and a measurable subset E of $\mathbb{R}^l$, we have*

$$\mathbb{P}[M(x) \in E] \leq e^{\epsilon}\mathbb{P}[M(x') \in E] + \delta$$

The definition intuitively says that a randomized mapping M behaves similarly on similar input databases. That is, for any neighboring records, after an ϵ, δ-DP mechanism, their statistical behaviors are similar. Hence, it is difficult for any party to determine which record is the source of the given output. Here, the privacy loss is given as

$$\mathcal{L} = \ln\left(\frac{\mathbb{P}[M(x) \in E]}{\mathbb{P}[M(x') \in E]}\right)$$

The Definition 2.1 ensures that for all neighbouring x, x' the absolute value of the privacy loss will be bounded by ϵ with probability at least $1 - \delta$. Smaller values of ϵ implies higher levels of privacy. When $\epsilon = \infty$, there is least privacy.

The key desirable properties of Differential Privacy are:

- Protection of arbitrary risks
- Automatic neutralization of linkage attacks
- Quantification of privacy loss
- Composition
- Group Privacy
- Closure under post-processing

They key idea behind Differential Privacy is that a user is given plausible deniability by adding random noise to their input. In the centralized Differential Privacy setting noise is added to the database. As the type of noise added is known statistical queries can still be computed by filtering out the noise while maintaining each user's individual privacy. However, this approach to Differential Privacy requires users to have trust the database maintainer to keep their privacy. To overcome this, a stronger privacy guarantee for users can be given by local setting where there is no need to trust a centralized authority. This is called *Local Differential Privacy.*

2.1 Local Differential Privacy (LDP)

Kasivishwanathan et.al [22] were first to formalize Local Differential Privacy. Unlike DP, in the LDP setting there is no trusted centroid curator. In this setting each user encodes their own data while inputing before it is transmitted to the untrusted server. This server can then compute statistical queries on the input data. The local settings poses its own challenges, as each user perturbs their input individually the overall variance is high. This has been a quickly developing field with many new state-of-the-art approaches to important problems in recent years [6, 30].

LDP is practically deployed at major organisations like Apple [9], to collect usage statistics and find commonly used emojis, new words that are not part of the dictionary

[1] Note that the summaries of the papers presented here are our understanding; if any error found, we are open to revisit.

yet and to improve user behaviour. Google [Read1] uses it in Chrome to collect commonly chosen homepages, settings, and other web browsing behaviour. Microsoft [10], uses it for their collection of telemetry data.

As mentioned before, Local Differential Privacy (LDP) gives users stronger privacy guarantees compared to the centralized notion of Differential Privacy where users still have to trust a central authority to keep their privacy. In LDP two basic types of protocols are often used: interactive and non-interactive. Both projects from Google and Apple use the non-interactive model. Moreover, implementing efficient interactive protocols in such applications is more difficult due to the latency of the network.

3

Multi-Armed Bandits (MAB)

The Multi-Armed Bandit (MAB) [27] problem provides a classic model for abstracting sequential decision making under uncertainty, and has attracted a wide range of interest in various areas, such as communication networks, online advertising, clinical trials, product testing, etc. In an MAB model, there is a set of arms, and each pull of an arm generates a random reward according to some unknown latent distribution of this arm. The agent adaptively chooses arms to pull according to past observations in order to achieve some goal. A widely studied goal is regret minimization, where the regret is the expected gap between a proposed algorithm and an optimal algorithm that knows the latent distributions. To minimize the regret, the agent needs to balance the trade-off between exploration and exploitation, where exploration refers to learning the environment and exploitation refers to pulling the best arm according to the current knowledge.

There are three fundamental formalizations of the bandit problem depending on the nature of the reward process: stochastic, adversarial, and Markovian. Three distinct strategies have been shown to effectively address each specific bandit model: the UCB algorithm [3] in the stochastic case, the Exp3 randomized algorithm in the adversarial case, and the so-called Gittins indices [18] in the Markovian case.

Most of the MAB problems studied till now are Non-Private MAB problems. For non-private bandit problems, either frequentist methods like UCB (Upper Confidence Bound) [4] or Bayesian methods like Thompson Sampling [2] have been shown to achieve optimal regret performance (up to constant factors). Recent days MAB settings are studied along with privacy guarantees. We discuss more about Private MAB problems in the following section.

3.1 Privacy in MAB Setting

In stochastic bandit systems, the rewards may refer to the users' activities, which may involve private information and the users may not want the agent to know. However, in many cases, the agent needs to know these activities to provide better services such as recommendations and news feeds. To handle this dilemma, Differential Privacy is adopted in to MAB algorithms to study regret bounds with a given privacy guarantee.

Here, we take clinical trials as an example to illustrate the use of DP in MAB. In an experiment of an illness with multiple treatments (referred as, arms), the experimenter (referred as, agent) wants to sequentially choose treatments for patients (referred as users) based on past observations on treatment effects. This problem can be viewed as an MAB regret minimization problem. However, the patients may not be willing to share the actual effects of the treatments with the experimenter due to privacy concerns. By the DP bandit algorithms, the actual effects will not be known by the experimenter, which provides a certain level of privacy guarantee to all patients, while also enabling the experimenter to learn from the observations efficiently. There are many DP bandit problems that are studied previously in literature [19, 26, 28].

Recently researchers are actively exploring LDP bandit problems. An example of LDP MAB problem is shopping websites. Here, the server wants to sequentially choose products to recommend according to the users' past purchase histories. However, some users are not willing to share this information as the purchase histories may reveal private information (e.g., a person who buys a lot of heart-disease medicines is more likely to have related illness). In this scenario, it is unlikely that there is a third-party centroid curator that can gain access to all the purchase data, since these data are commonly viewed as a valuable property for the company's business success. This indicates the necessity of the LDP setting instead of the DP setting.

In the remainder of the report, we summarize a few papers related to the concepts discussed above.

4

RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response

[Read1]
Author: Úlfar Erlingsson, Vasyl Pihur, Aleksandra Korolova
Appeared in: ACM SIGSAC- CCS 2014

4.1 Problem Addressed

RAPPOR is a technology for crowdsourcing statistics from end-user, anonymously, with strong privacy guarantees. RAPPOR permits data collection with strong privacy guarantees for each client, and without linkability of their reports, while assuring high utility analysis of the collected data. This paper describes and motivates RAPPOR, details its differential-privacy and utility guarantees.

4.2 Previous Work

- ϵ-Differential Privacy guarantee [13] protects the privacy of specific respondent, but the privacy degrades if the survey is conducted repeatedly.
- Randomized Response, a survey technique [31] was introduced for collecting statistics on sensitive topics where respondents wish to retain confidentially.

4.3 Novelty

- RAPPOR provides longitudinal privacy, i.e., protects the privacy of clients even if the data is collected repeatedly from them (even infinitely often).
- RAPPOR collects statistics about arbitrary set of strings by applying randomized response to Bloom filters.
- RAPPOR does not require any trusted third party, it is performed locally on the client side.
- RAPPOR provides high-utility decoding framework for learning statistics based on a sophisticated combination of hypotheses testing, least-squares solving, and LASSO regression.

4.4 RAPPOR

RAPPOR is built on the basic idea of memoization and provides framework for one-time and longitudinal privacy by playing the randomized game twice and with a memoization step in between. The RAPPOR algorithm from the client-side reports a bit array of size k (execution parameter) that encodes noisy representation of value v. This reveals a controlled amount of information about v, limiting the server's ability to learn v with confidence.

The first randomization is, *permanent randomized response* used to create a noisy answer which is memoized by the client and permanently reused in place of real answer. This step ensures longitudinal privacy of the client. The second is, *instantaneous randomized response* protects data against possible attacks from externalities.

In first step, memoization is important, as only randomization is not sufficient to maintain privacy in the face of multiple collections. Memoization is effective when the true value of the client does not change or change in uncorrelated fashion.

4.4.1 Randomization Algorithm

RAPPOR takes in client's true value v and execution parameters k,h,f,p,q and is executed locally on the client's machine performing the following steps:

1. **Signal.** Hash clients value v on to bloom filter of size k using h hash functions.

2. **Permanent randomized response**. For each client's value v and bit i, $0 \leq i < k$ in B, create a binary reporting value B'_i which is defined as follows:

$$B'_i = \begin{cases} 1, \text{with probability } \frac{1}{2}f \\ 0, \text{with probability } \frac{1}{2}f \\ B_i, \text{with probability } 1-f \end{cases}$$

   where f is a parameter controlling level of longitudinal privacy guarantee.

   This value B'_i is memoised and reused in future for the distinct value v.

3. **Instantaneous randomized response**. In a bit array S of size k and initialize to 0. Each bit is set to 1 with probability

$$P(S_i = 1) = \begin{cases} q, \text{if } B'_i = 1 \\ p, \text{if } B'_i = 0 \end{cases} \tag{1}$$

4. The generated report S is sent to server.

Theorem 4.1. *The permanent randomized response satisfies ϵ_∞-Differential Privacy where* $\epsilon_\infty = 2h \ln\left(\frac{1-\frac{1}{2}f}{\frac{1}{2}f}\right)$

Theorem 4.2. *The instantaneous randomized response satisfies ϵ_1-Differential Privacy where* $\epsilon_1 = h\ln\left(\frac{q^*(1-p^*)}{p^*(1-q^*)}\right)$

where $p^* = P(S_i = 1|b_i = 1) = \frac{1}{2}f(p+q) + (1-f)q$, $q^* = P(S_i = 1|b_i = 0) = \frac{1}{2}f(p+q) + (1-f)p$

4.4.2 Utility Decoding

Each client in RAPPOR is randomly assigned and becomes a permanent member of one of the m cohorts. Each cohort implements different set of hash functions. The approach adapted for learning is as follows:

- Estimate the number of times each bit i within cohort j, t_{ij} in set B for each cohort. Given the number of times each bit i in cohort j, c_{ij} was set in a set of N_j reports,

$$t_{ij} = \frac{c_{ij} - (p + \frac{1}{2}fq - \frac{1}{2}fp)N_j}{(1-f)(q-p)}$$

- Let Y be the matrix of values t_{ij}, $i \in [1,k], j \in [1,m]$
- Create a matrix X of size $km \times M$ where M is the number of candidate strings under consideration. X is a matrix with 1's at bloom filter bits for each string for each cohort. Otherwise 0.
- Use a Lasso regression model to fit a model $Y \approx X$ and select candidate strings corresponding to nonzero coefficients.
- Fit a regular least-squares regression using the selected variables to estimate counts, their standard errors and p-values.
- Compare p-values to a Bonferroni corrected level of $\frac{\alpha}{M} = \frac{0.05}{M}$ to determine which frequencies are statistically significant from 0. Alternatively, controlling the False Discovery Rate (FDR) at level α using the Benjamini-Hochberg procedure.

5

Differentially Private and Fair Deep Learning: A Lagrangian Dual Approach

[Read4]
Author: Cuong Tran, Ferdinando Fioretto, Pascal Van Hentenryck
Appeared in: Arxiv

5.1 Problem Addressed

A critical concern in data-driven decision making is to build models whose outcomes do not discriminate against some demographic groups. To ensure non-discrimination in learning tasks, knowledge of the sensitive attributes is essential, while, in practice, these attributes may not be available due to legal and ethical requirements. To address this challenge, this paper studies a model that protects the privacy of the individuals sensitive information while also allowing it to learn non-discriminatory predictors.

5.2 Previous Work

- Abadi et.al [1] derived Differential Privacy version of stochastic gradient descent (DP-SGD) and proposed technique *moment accountant* to track detailed information of privacy loss incurred by a sequence of SGD steps.
- Other type of work avoids using the iterative nature of the optimization algorithms, by exploiting a collection of teacher models [8] to train a privacy-preserving model.
- [12], one of the earliest contribution linking fairness and Differential Privacy. Shows that individual fairness is a generalization of Differential Privacy.
- [23] has observed that private models may have a negative impact towards fairness.

5.3 Problem Setting

Let D be dataset consisting of n individual data points (X_i, A_i, Y_i) with $i \in [n]$ drawn i.i.d from an unknown distribution. $X_i \in \mathcal{X}$ is a non-sensitive feature vector, $A_i \in \mathcal{A} = [m]$ is a protected attribute, and $Y_i \in \mathcal{Y} = \{0,1\}$ is a binary label. The goal is to learn a classifier $\mathcal{M}_\theta : \mathcal{X} \leftarrow \mathcal{Y}$ where θ is vector of real-valued parameters, that ensures a specified non-discriminatory notion with A while guaranteeing the privacy of sensitive attribute A. The quality of the model is measured in terms of non-negative, differentiable loss function $\mathcal{L} : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}_+$. The problem minimizes empirical risk function:

$$\min_\theta J(\mathcal{M}_\theta, D) = \frac{1}{n}\sum_{i=1}^{n} \mathcal{L}(\mathcal{M}_\theta(X_i), Y_i)$$

5.4 Approach

This paper presents a Lagrangian dual method to enforce several fairness constraints directly into the training cycle of a deep neural network and proposes a differentially private and fair version of the learning algorithm. We detail this in the following sections.

5.5 Fairness

The fairness constraints considered are focused on the three fairness notions, 1. Demographic Parity, 2. Equalized Odds, 3. Accuracy Parity. These constraints are relaxed into objectives using Lagrangian relaxation. This relaxation relies on an iterative scheme that interleaves the learning of a number of Lagrangian relaxations with a subgradient method to learn the best multipliers. This method is called Fair-Lagrangian Dual (F-LD).

5.6 Privacy

The computation of these gradients can be made differentially private by the introduction of carefully calibrated Gaussian noise. This is performing a Differentially Private Stochastic Gradient Descent (DP-SDG) step. DP-SDG computes the gradients for each data sample in a random mini-batch, clips their L2-norm, computes the average, and adds noise to ensure privacy. This extension to F-LD guarantees both fairness and privacy, it is called Private and Fair Lagrangian Dual (PF-LD).

The privacy analysis of PF-LD relies on the moment accountant for Sampled Gaussian (SG) mechanism, whose privacy is analyzed using Rényi Differential Privacy (RDP).

Theorem 5.1. *For a given function $f : D \to \mathbb{R}^d$, such that $||f(D) - f(D')||_2 \leq 1$ for any neighbouring databases D, D' the SG mechanism $SG_{q,\sigma}$ with sampling ratio q and standard deviation σ satisfies (α, ϵ)-RDP with:*

$$\epsilon \leq \mathcal{D}_\alpha[\mathcal{N}(0,\sigma^2)||(1-q)\mathcal{N}(0,\sigma^2) + q\mathcal{N}(1,\sigma^2)]$$
$$\epsilon \leq \mathcal{D}_\alpha[(1-q)\mathcal{N}(0,\sigma^2) + q\mathcal{N}(1,\sigma^2)||\mathcal{N}(0,\sigma^2)]$$

Theorem 5.2. *PF-LD satisfies $(\alpha, \frac{T_{n\epsilon_p}}{|B|} + T\epsilon_d) - RDP$ where B is the mini batch at each iteration, T epochs and ϵ_p, ϵ_d are parameters satifying two equations in previous theorem, respectively.*

5.7 Conclusion

The paper proposed a framework to train deep learning models that satisfy several notions of group fairness, including equalized odds, accuracy parity, and demographic parity, while ensuring that the model satisfies Differential Privacy for the protected attributes. The framework relies on the use of Lagrangian duality to accommodate the fairness constraints and also showed how to inject calibrated noise to the primal and dual steps of the Lagrangian dual process to guarantee privacy of the sensitive attributes.

6

Multi-Armed Bandits with Local Differential Privacy

[Read3]
Author: Wenbo Ren, Xingyu Zhou, Jia Liu, Ness B. Shroff
Appeared in: Arxiv

6.1 Problem Addressed

This paper focuses at regret minimization for Multi-Armed Bandit (MAB) problems with Local Differential Privacy guarantee (LDP). In stochastic bandit systems, the rewards may refer to the users' activities, which may involve private information and the users may not want the agent to know. For this, paper adopts Differential Privacy and study the regret upper and lower bounds for MAB algorithms with a given LDP guarantee.

6.2 Model Overview

6.2.1 Bandit Model

The bandit model has n arms indexed by $1, 2, 3, \dots, n$, each arm is associated with unknown latent distribution. Every t-th pull of arm a returns a reward R_a^t according to its distribution. $\mu_a = \mathbb{E}[R_a^t]$ is the mean reward of arm a. The maximum mean reward among the arms is denoted as, $\mu^* = \max_{a \in n} \mu_a$. With this, for every arm a gap is defined as, $\Delta_a = \mu^* - \mu_a$. An arm is said to be optimal if $\Delta_a = 0$, and suboptimal otherwise.

6.2.2 Regret Minimisation

A^t denotes the t-th pulled arm and $N_a^t = \sum_{\tau=1}^t \mathbf{1}A^\tau = a$ denotes the number of pulls on arm a till time t. After T(time horizon) pulls the expected reward is $\sum_{t=1}^T \mathbb{E}[\mu_{A^t} = \mathbb{E}[\sum_{a\in[n]} N_a^T \mu_a]$ and regret is defined as

$$R(T) = T\mu^* - \sum_{t=1}^{T} \mathbb{E}[\mu_{A^t}] = \mathbb{E}[\sum_{a\in[n]} N_a^T \Delta_a]$$

In the paper, the LDP bandit model splits parties into three categories: agents, curators and users. Users don't trust the agents and thus, curators stand between the users and the agents, providing privacy to the users and minimizing regret for the agents. The agent makes a decision on which arm to pull according to the knowledge of past private responses and sends a request to a user's curator. The curator then 'pulls the arm', receives the reward, and returns a private response to the agent. This implies that agent does not know actual rewards.

Definition 6.1 (LDP Bandit Model)**.** *Let A^t be the t-th pulled arm, R^t be the corresponding reward. For $\epsilon > 0$, the bandit model is said to be ϵ-LDP if there is an ϵ-DP mechanism $M : \mathcal{D} \to \mathbb{R}$ and $A^{t+1} \in \sigma(A^s, M(R^s) : 1 \leq s \leq t)$ for any time t.*

6.3 Previous Work

- Gazane et.al [16] proposed an LDP bandit algorithm that works for arms with Bernoulli rewards.

- Basu et.al [7] studied LDP bandit problem, in which distribution-dependent and distribution-free regret lower bounds were proved.

6.4 Lower Bound

Theorem 6.1 (Lower Bound)**.** *Let $\epsilon > 0$ be given. Assume that the rewards of all the arms follow Bernoulli distributions. The regret $R(T)$ of any ϵ-LDP policy satisfies*

$$\lim_{T\to\infty} \frac{R(T)}{\log(T)} \geq \frac{1}{(e^{\epsilon} - e^{-\epsilon})^2} \sum_{a:\Delta_a > 0} \frac{1}{\Delta_a}$$

When $\epsilon \to 0$, since $e^{\epsilon} - e^{-\epsilon} \simeq 2\epsilon$, we have $\lim_{T\to\infty} \frac{R(T)}{\log(T)} \geq \frac{1}{4\epsilon^2} \sum_{a:\Delta_a>0} \frac{1}{\Delta_a}$

6.5 Algorithms and Upper Bounds

This paper proposes two ϵ-DP mechanisms: one converts bounded rewards to Laplace responses and other converts to Bernoulli responses. Later agents use these private responses to trade off the exploration and exploitation with UCB like methods.

Mechanism CTL(ϵ) at the curator, adds laplacian noise to each reward while preserving the ϵ-DP. UCB algorithm associated with this, LDP-UCB-L(ϵ), is ϵ-LDP and its distribution-free regret is at most $O(\epsilon^{-1}\sqrt{nT\log T})$.

Algorithm 1 LDP UCB algortihm with Laplacian mechanism

```
Pull each arm and receive the private response.
Define $\hat{\mu}_a^t$, empirical mean of private responses of arm $a$ till time $t$
Define $N_a^t$, number of pulls of arm $a$ and $\mu_a^t = \hat{\mu}_a^t + \sqrt{\frac{(2\log t)}{N_a^t}} + \sqrt{\frac{(32\log t)}{\epsilon^2 N_a^t}}$
while $t < T$ do
    if there is an arm $a$ with $N_a^t \leq 4\log(t+1)$ then : $a^t \leftarrow a$
    else : $a^t \leftarrow \mathrm{argmax}_{a\in[n]}\, u_a^t$
    end if
    Pull arm $a^t$ and receive private response.
    $t \leftarrow t+1$; Update $\hat{\mu}_a^t$, $N_a^t$ and $u_a^t$
end while
```

Similarly, Mechanism CTB(ϵ) adds Bernoulli noise to each reward. The UCB algorithm associated with this, LDP-UCB-B(ϵ) is ϵ-LDP and its distribution-free regret is at most $O(\epsilon^{-1}\sqrt{nT\log T})$.

Algorithm 2 LDP UCB algortihm with Bernoulli mechanism

```
Pull each arm and receive the private response.
Define $\hat{\mu}_a^t$, empirical mean of private responses of arm $a$ till time $t$
Define $N_a^t$, number of pulls of arm $a$ and $\mu_a^t = \hat{\mu}_a^t + \sqrt{\frac{(2\log t)}{N_a^t}}$
while $t < T$ do
    $a^t \leftarrow \mathrm{argmax}_{a\in[n]}\, u_a^t$
    end if
    Pull arm $a^t$ and receive private response.
    $t \leftarrow t+1$; Update $\hat{\mu}_a^t$, $N_a^t$ and $u_a^t$
end while
```

For unbounded rewards, the rewards are mapped to sigmoid function $s(r) = (1+e^{-r})^{-1}$ with this outputs are guaranteed to be in [0,1].

7

Local differential privacy for social network publishing

[Read2]
Author: Peng Liu, YuanXin Xu, Quan Jiang, Yuwei Tang, Yameng Guo, Li-e Wang, Xianxian Li
Appeared in: NeuCom 2020

7.1 Problem Addressed

Social networks often contain sensitive information, thus, we need to protect privacy when publishing social network data. However, the current differential privacy for social network data publishing seriously influences the structure of the social network. The paper proposes new local differential privacy method that can offer a better way to protect a network's structural information.

The goal of the paper is to determine how to publish local differential privacy social network and minimize the influence of noise on the structure information of the data.

7.2 Previous Work

- Many studies have investigated anonymized techniques [20, 21] to ensure network data privacy.

- Anonymized models cannot resist several newly discovered privacy attacks and still lead to privacy breaches [25].

- Differential privacy [11] has been proposed to solve this vulnerability.

7.3 Local Differential Privacy Model

Definition 7.1 (Community in Social Network)**.** *Let $G = (V, E)$ be a social network, the set of communities $C = \{C_1, C_2, \dots, C_m\}$ where $C_i \cap C_j \neq \emptyset \quad \forall 1 \leq i \neq j \leq m$. For each $v \in C$, if the density of internal connection is higher than outside, then the elements in C is a community.*

This paper uses an efficient community detection algorithm, fast-Newman algorithm, to discover community structure and a modularity increase optimization method to explore the community structure. The modularity function is defined as

$$Q = \sum_{c=1}^{m} \left[\frac{l_c}{n} - (\frac{d_c}{2n})^2 \right]$$

where m is the number of communities, l_c is total number of edges and d_c is the sum of the degrees of nodes in community c, and n is number of nodes in G.

The goal is to release a sanitized network $\tilde{G}$ satisfying ϵ-differential privacy while trying to maintain the original structural information of the data. The method includes two processes: injecting noise into the community and creating disturbances between communities.

The reconstruction of edges in the community is defined according to the classic generation graph model

Definition 7.2 (Probabilistic Edge Reconstruction)**.** *The probability of an edge is proportion to the degree of the two connected nodes. The probability of an edge P_{c_i} is defined as:*

$$P_{c_i} = (E_{c_i,c_j} | \theta_{c_i}) = \frac{\theta_{d_i} \theta_{d_j}}{\sum_{v_k \in V} \theta_{d_k}}$$

where $\theta_{c_i} = [\theta_{d_1}, \dots, \theta_{d_{N_V}}]$ is model learning the node degree from the original social network community, and $\theta_{d_{c_i}} = d_{c_i}$.

The probabilistic edge reconstruction in a community ensures that the generated graph is equal to the expected degree of the original graph.

7.3.1 Algorithm

The algorithm given in the paper consists of three steps:

1. Community Detection

2. Laplace noise is injected into a community

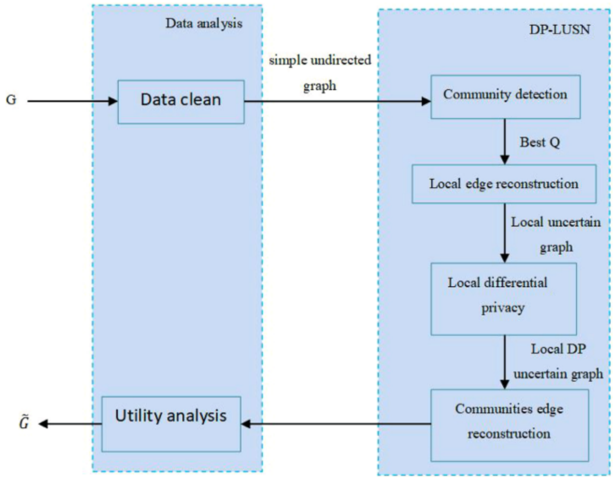

Figure 1: DP-LUSN algorithm

3. Injecting noise in a community and creating disturbances outside a community

Fast-Newman algorithm is used for community partitioning, i.e., Step 1. Step 2 of the algorithm achieves local differential privacy in a social network via uncertain graph (DP-LUSN). The DP-LUSN algorithm satisfies ϵ-differential privacy. In final step, algorithm generates probability edges between communities.

8

Principal Component Analysis in the Local Differential Privacy Model

[Read5]
Author: Di Wang, Jinhui Xu
Appeared in: IJCAI 2019

8.1 Problem Addressed

This paper studies Principal Component Analysis (PCA) under the non-interactive local differential privacy model and aims to find the limitations and optimal algorithm of PCA under the non-interactive local differential privacy model.

PCA is a fundamental technique for dimension reduction in statistics. Big data now ubiquitously exist in our daily life, which need to be analyzed (or learned) statistically by methods like regression and PCA. However, due to the presence of sensitive data and their distributed nature, such data are extremely difficult to aggregate and learn from. A better solution is to use some differentially private mechanisms to conduct the aggregation and learning tasks.

8.2 Approach

For the low dimensional case, paper gives the optimal ratefor the private minimax risk of the k-dimensional PCA using the squared subspace distance as the measurement. For the high dimensional row sparse case, it first gives a lower bound on the private minimax risk and then provide an efficient algorithm to achieve a near optimal upper bound.

8.3 Previous Work

- For the low dimensional case, [5] studied the private PCA problem under the interactive local differential privacy model and introduced an approach based on the noisy power method.
- For the private high dimensional sparse PCA, [17] proposed a noisy iterative hard thresholding power method, which is an interactive LDP algorithm.
- This paper adopts the optimal procedure based on perturbing the covariance by Gaussian matrices, which has been studied in [15].

8.4 Minimax Risk

8.4.1 Classical Minimax Risk

Let $\mathcal{P}$ be a class of distributions over a data universe χ. For each distribution $P \in \mathcal{P}$, there is a deterministic function $\theta(P) \in \Theta$, where Θ is the parameter space. Let $\rho : \Theta \times \Theta :\to R_+$ be a semi-metric function on space Θ and $\Phi : R_+ \to R_+$ be a non-decreasing function with $\Phi(0) = 0$.

The minimum risk in metric Φ_{op} is defined by the following saddle point problem:

$$\mathcal{M}_n(\theta(\mathcal{P}), \Phi_{op}) = \inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_p[\Phi(\rho(\hat{\theta}(X_1, \ldots, X_n), \theta(P)))]$$

8.4.2 Private Minimax Risk

Given a family of distributions $\theta(P)$ and privacy parameter $\epsilon > 0$, the ϵ non-interactive private minimax risk in the metric Φ_{op} is

$$\mathcal{M}_n^{Nint}(\theta(P), \Phi_{op}, \epsilon) = \inf_{Q \in Q_\epsilon} \mathcal{M}_n(\theta(P), \theta_{op}, Q)$$

Here Q_ϵ is the set of all ϵ non-interactively locally differentially private mechanisms.

8.5 Results

Let S and S' be two k-dimensional subspaces in $\mathbb{R}^p$ and $E = VV^T$, $F = WW^T$ be orthogonal matrices corresponding to them. Then the squared subspace distance between S and S' is defined by

$$|| \sin \Theta(S.S')||_F^2 = ||E - F||_F^2 = \frac{1}{2}||VV^T - WW^T||_F^2$$

The paper proves that the minimax risk (measured by the squared subspace distance) under ϵ non-interactive local differential privacy (LDP) is lower bounded by $\Omega(\frac{\lambda_1 \lambda_{k+1} pk}{(\lambda_k - \lambda_{k+1})^2 n \epsilon^2})$, where p is the dimensionality of the data and n is the number of data records. It also shows the $\Omega(\frac{pk}{n\epsilon^2})$ is optimal.

As we observe that the above result is too large in high dimensions. Thus, for high dimensional case, paper considers the sparse PCA under non-interactive local model and proves that the private minimax risk is lower bounded by $\Omega(\frac{\lambda_1 \lambda_{k+1} ks \log p}{(\lambda_k - \lambda_{k+1})^2 n \epsilon^2})$, where λ_1,λ_k and λ_{k+1} are eigenvalues and s is the sparsity parameter of the eigen vectors. It also gives an optimal upper bound of $O(\frac{\lambda_1^2 s^2 \log p}{(\lambda_k - \lambda_{k+1})^2 n \epsilon^2})$

9

Overview

Paper	**Published in**	**Problem Statement**	**Solution Proposed**
Differential Privacy			
Differentially Private and Fair Deep Learning: A Lagrangian Dual Approach	arXiv 2009	To ensure non-discrimination in learning tasks, knowledge of the sensitive attributes is essential, while, in practice, these attributes may not be available. The paper focuses on learning general classifiers, such as neural networks, that satisfy group fairness and protect the disclosure of the sensitive attributes.	The method relies on the notion of Differential Privacy and the use of Lagrangian duality to design neural networks that can accommodate fairness constraints while guaranteeing the privacy of sensitive attributes.
A bounded-noise mechanism for Differential Privacy	arXiv 2020	Paper studies answering multiple counting queries	Present an (ϵ, δ)-private mechanism with optimal error
Applications of Local Differential Privacy			
RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response	ACM SIGSAC-CCS 2014	This paper studies about RAPPOR, a privacy preserving platform being used by Google	Paper details RAPPOR's differential-privacy and utility guarantees, discusses its practical deployment and properties in the face of different attack models
Local Differential Privacy for social network publishing	NeuCom 2020	the current Differential Privacy for social network data publishing seriously influences the structure of the social network	The paper proposes Local Differential Privacy model for social network publishing that preserves community structure information
Principal Component Analysis in the Local Differential Privacy Model	IJCAI 2019	This paper studies PCA under the non-interactive Local Differential Privacy model	The paper provides a minimax risk lower bound under ϵ non-interactive Local Differential Privacy for both low and high dimensional settings.
Local Differential Privacy in MAB Setting			
Multi-Armed Bandits with Local Differential Privacy	arXiv 2007	This paper studies the Multi-Armed Bandit problem with Local Differential Privacy guarantee	Provides a tight regret lower bound and proposes algorithms with tight regret upper bounds

Table 1: **Research Papers Summarised**

Acknowledgement

First author would like to thank Mr. Sankarshan Damle for useful discussions.

List of Papers Summarized

[Read1] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, CCS '14, page 1054–1067, New York, NY, USA, 2014. Association for Computing Machinery.

[Read2] Peng Liu, YuanXin Xu, Quan Jiang, Yuwei Tang, Yameng Guo, Li-e Wang, and Xianxian Li. Local differential privacy for social network publishing. *Neurocomputing*, 391:273–279, 2020.

[Read3] Wenbo Ren, Xingyu Zhou, Jia Liu, and Ness B. Shroff. Multi-armed bandits with local differential privacy, 2020.

[Read4] Cuong Tran, Ferdinando Fioretto, and Pascal Van Hentenryck. Differentially private and fair deep learning: A lagrangian dual approach, 2020.

[Read5] Di Wang and Jinhui Xu. Principal component analysis in the local differential privacy model. pages 4795–4801, 08 2019.

References

[1] Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, Oct 2016.

[2] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem, 2012.

[3] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 05 2002.

[4] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 05 2002.

[5] Maria Florina Balcan, Simon S. Du, Yining Wang, and Adams Wei Yu. An improved gap-dependency analysis of the noisy power method, 2016.

[6] Raef Bassily, Kobbi Nissim, Uri Stemmer, and Abhradeep Thakurta. Practical locally private heavy hitters, 2017.

[7] Debabrota Basu, Christos Dimitrakakis, and Aristide Tossou. Differential privacy for multi-armed bandits: What is it and what is its cost?, 2020.

[8] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin Raffel. Mixmatch: A holistic approach to semi-supervised learning, 2019.

[9] Graham Cormode, Somesh Jha, Tejas Kulkarni, Ninghui Li, Divesh Srivastava, and Tianhao Wang. Privacy at scale: Local differential privacy in practice. In *Proceedings of the 2018 International Conference on Management of Data*, SIGMOD '18, page 1655–1658, New York, NY, USA, 2018. Association for Computing Machinery.

[10] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately, 2017.

[11] C Dwork and NISSIM K MCSHERRYF. Calibratingnoiseto sensitivityinprivatedataanalysis. *ProceedingsofThird TheoryofCryptographyConference, March4G7*, 2006.

[12] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Rich Zemel. Fairness through awareness, 2011.

[13] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. volume Vol. 3876, pages 265–284, 01 2006.

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

[15] Cynthia Dwork, Kunal Talwar, Abhradeep Thakurta, and Li Zhang. Analyze gauss: Optimal bounds for privacy-preserving principal component analysis. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '14, page 11–20, New York, NY, USA, 2014. Association for Computing Machinery.

[16] Pratik Gajane, Tanguy Urvoy, and Emilie Kaufmann. Corrupt bandits for preserving local privacy, 2017.

[17] J. Ge, Zhaoran Wang, Mengdi Wang, and Han Liu. Minimax-optimal privacy-preserving sparse pca in distributed systems. In *AISTATS*, 2018.

[18] John Gittins, Kevin Glazebrook, and Richard Weber. *Multi-Armed Bandit Allocation Indices, 2nd Edition*, volume 33. 02 2011.

[19] Awni Hannun, Brian Knott, Shubho Sengupta, and Laurens van der Maaten. Privacy-preserving multi-party contextual bandits, 2020.

[20] Michael Hay, Gerome Miklau, David Jensen, Don Towsley, and Philipp Weis. Resisting structural re-identification in anonymized social networks. *Proc. VLDB Endow.*, 1(1):102–114, August 2008.

[21] Jia Jiao, Peng Liu, and Xianxian Li. A personalized privacy preserving method for publishing social network data. In *International Conference on Theory and Applications of Models of Computation*, pages 141–157. Springer, 2014.

[22] Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately?, 2010.

[23] Satya Kuppam, Ryan Mckenna, David Pujol, Michael Hay, Ashwin Machanavajjhala, and Gerome Miklau. Fair decision making using privacy-protected data, 2020.

[24] Noman Mohammed, Rui Chen, Benjamin C.M. Fung, and Philip S. Yu. Differentially private data release for data mining. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '11, page 493–501, New York, NY, USA, 2011. Association for Computing Machinery.

[25] A. Narayanan and V. Shmatikov. De-anonymizing social networks. In *2009 30th IEEE Symposium on Security and Privacy*, pages 173–187, 2009.

[26] Roshan Shariff and Or Sheffet. Differentially private contextual linear bandits, 2018.

[27] D. Stoyan. Berry, d. a. and b. fristedt: Bandit problems. sequential allocation of experiments. monographs on statistics and applied probability. chapman and hall, london/new york 1985, 275 s. *Biometrical Journal*, 29(1):20–20, 1987.

[28] Aristide C. Y. Tossou and Christos Dimitrakakis. Algorithms for differentially private multi-armed bandits. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16, page 2087–2093. AAAI Press, 2016.

[29] Baoxiang Wang and Nidhi Hegde. Privacy-preserving q-learning with functional noise in continuous state spaces, 2019.

[30] T. Wang, N. Li, and S. Jha. Locally differentially private frequent itemset mining. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 127–143, 2018.

[31] Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.