Towards Practical Differential Privacy in Data Analysis: Understanding the Effect of Epsilon on Utility in Private ERM

Yuzhe Li
Institute of Information Engineering
Chinese Academy of Sciences
liyuzhe@iie.ac.cn

Yong Liu[*]
Gaoling School of Artificial Intelligence
Renmin University of China, Beijing, China
liuyonggsai@ruc.edu.cn

Bo Li
Institute of Information Engineering
Chinese Academy of Sciences
libo@iie.ac.cn

Weiping Wang
Institute of Information Engineering
Chinese Academy of Sciences
wangweiping@iie.ac.cn

Nan Liu
Institute of Information Engineering
Chinese Academy of Sciences
liunan@iie.ac.cn

ABSTRACT

As computation over sensitive data has been an important goal in recent years, privacy-preserving data analysis has gradually attracted more and more attention. Among various mechanisms, differential privacy has been widely studied due to its formal privacy guarantees for data analysis. As one of the most important issues, the crucial trade-off between the strength of privacy guarantee and the effect of analysis accuracy is highly concerned among the researchers. Existing theories for this issue consider that the analyst should first chooses a privacy requirement and then attempts to maximize the utility. However, as differential privacy is gradually deployed in practice, a gap between theory and practice comes out: in practice, product requirements often impose hard accuracy constraints, and privacy (while desirable) may not be the over-riding concern. Thus, it is usually that the requirement of privacy guarantee is adjusted according to the utility expectation, not the other way around. This gap raises the question of how to provide maximum privacy guarantee for data analysis due to a given accuracy requirement. In this paper, we focus our attention on private Empirical Risk Minimization (ERM), which is one of the most commonly used data analysis method. We take the first step towards solving the above problem by theoretically exploring the effect of ϵ (the parameter of differential privacy that determines the strength of privacy guarantee) on utility of the learning model. We trace the change of utility with modification of ϵ and reveal an established relationship between ϵ and utility. We then formalize this relationship and propose a practical approach for estimating the utility under an arbitrary value of ϵ. Both theoretical analysis and experimental results demonstrate high estimation accuracy and broad applicability of our approach in practical applications. As providing algorithms with strong utility guarantees that also give privacy when possible becomes more and more accepted, our approach would have high practical value and may be likely to be adopted by companies and organizations that would like to preserve privacy but are unwilling to compromise on utility.

Keywords differential privacy · machine learning · parameter selection

---

[*]Corresponding author.

1 Introduction

In recent years, more and more researches have exposed potential privacy risks in data analysis tasks Fredrikson et al. [2015], Tramèr et al. [2016], Shokri et al. [2017], Ganju et al. [2018], Yeom et al. [2018], Melis et al. [2019]. Therefore, with more and more sensitive data is involved in data analysis tasks, the privacy concerns are becoming more and more serious. To solve this problem, various mechanisms have been proposed to preserve privacy leakage during data analysis, e.g., *k*-anonymity and its variants Sweeney [2002], Hay et al. [2008], Liu and Terzi [2008], Li et al. [2012]. However, due to the limitations of the schemes (e.g., they are syntactic properties based) and the rich amount of auxiliary information available to adversaries, these mechanisms are susceptible to emerging *structure based de-anonymization attacks* Backstrom et al. [2011], Narayanan and Shmatikov [2009], Srivatsa and Hicks [2012]. One definition that overcome the limitations of existing mechanisms by protecting against any adversaries with any potential auxiliary information is differential privacy.

Differential privacy Dwork et al. [2006] provides a generic solution for privacy-preserving data analysis. It proves formal guarantees, in terms of a privacy budget ϵ, on how much information is leaked. However, since approaches for achieving differential privacy necessarily lead to a decreased accuracy, it is necessary for differentially private data analysis approaches to simultaneously consider privacy guarantees and utility guarantees. Privacy alone can be easily achieved by injecting adequate noise, but this usually leads to useless analysis results. Mathematically, the trade-off between privacy and utility is tunable by ϵ, that is, smaller ϵ leads to stronger privacy guarantee with lower utility, and larger ϵ leads to weaker privacy guarantee with higher utility. Thus, the choice of ϵ really matters for the implementation of differential privacy in data mining.

Traditional approaches to differential privacy assume that the user can pick the privacy parameter ϵ via some exogenous process, and attempt to maximize the utility of the computation subject to the privacy constraint Chaudhuri et al. [2011], Kifer et al. [2012], Bassily et al. [2014], Wang et al. [2017], Iyengar et al. [2019]. But these privacy-first approaches often encounter problems in practical applications. The first problem would be that they place an inordinate burden on the data analyst to understand differential privacy. These approaches imply a precondition that the user is capable to perceive and judge the effects of a privacy parameter. However, as the definition of differential privacy offers little insight into how this should be done, it may be beyond the capabilities of users. The second problem is that these approaches do not provide any guarantees to the users on the quality they really care about, namely the utility. Due to the lack of available theories, it is difficult for users to know in advance what impact the choice of ϵ will bring to the utility. The only theoretically sound method, namely "utility theorem" Chaudhuri et al. [2011], often fails to provide meaningful guidance as they tend to be worst-case bounds and would be too conservative as a practical reference Ligett et al. [2017]. The inability to assess the impact of utility increases the uncertainty of introducing privacy guarantee, thus may introduce unpredictable risks and costs. For example, low ϵ would lead to low accuracy, and low accuracy may lead to frequent errors in face recognition system, which leads to the introduction of human assistance.

To solve the problem of traditional methods, Ligett et al. [2017] considers a different perspective in the domain of private empirical risk minimization (ERM), where the goal is to give the strongest privacy guarantee possible, subject to a constraint on the acceptable accuracy. They propose a framework for developing differentially private algorithms under accuracy constraints, which allows one to choose a given level of accuracy first, and then finding the private algorithm meeting this accuracy. However, an obvious drawback of this method is that it requires so many attempts of ϵ (noise subtraction), that would take a lot of computation and training time. This is often difficult to be accepted in practice, especially in large-scale learning tasks.

In this paper, we continue the exploration of solving traditional problems from accuracy-first perspective in private ERM. We provide a novel practical method for assessing the impact of ϵ on utility of differentially private learning algorithms due to the inherent formulaic relations between them, which would be both as accurate as empirical approach and as efficient as theoretical methods. The main contributions of this paper can be summarized as follows:

- We comprehensively analyze the effect of ϵ on utility in objective-perturbation private ERM algorithms. We are the first to find out that there exists a data-dependent relationship between the value of ϵ and the utility of final trained learning model.

- We are the first to formalize this relationship and propose a practical approximate approach for utility analysis. We also show that our approximate approach could be easily used in practice applications.

- We conduct a series of evaluations over three real-world datasets. Experimental results show that the estimation results of our approach is considerable close to the actual measured results, which demonstrates high estimation accuracy and broad applicability of our approach in practical applications.
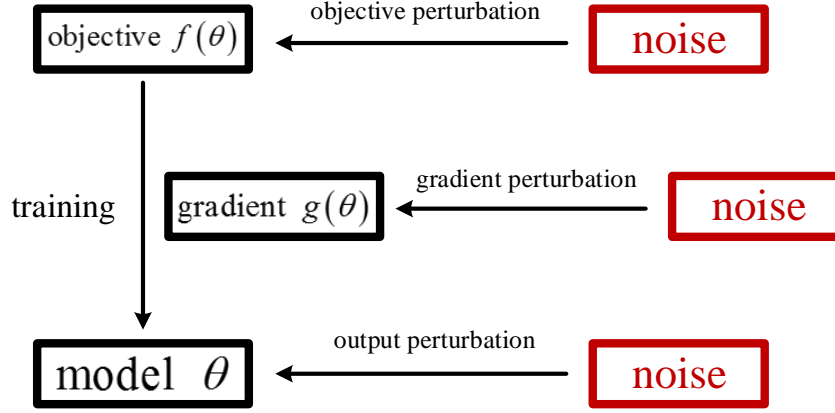
Figure 1: Privacy noise mechanisms for differentially private machine learning algorithm.

2 Preliminary

In the rest of this paper, we use $\| \cdot \|$ to denote the L_2 norm. All vectors will typically be written in boldface. We will use $\boldsymbol{I}$ to denote the identity matrix. Before delving into the details of our approach, we will first recall some basic concepts and expressions of differential privacy and differentially private learning.

2.1 Differential Privacy

Differential privacy is a rigorous mathematical framework for privacy guarantee. It has recently received a significant amount of research attention for its robustness to known attacks. The typical definition of differential privacy is given in Dwork et al. [2006], as follows:

Definition 1 ((ϵ, δ)-differential privacy). *Given a randomized function $\mathcal{M} : \mathcal{D} \to \mathcal{R}$ with domain $\mathcal{D}$ and range $\mathcal{R}$, we say $\mathcal{M}$ is (ϵ, δ)-differential privacy if for all pairs of neighboring inputs D, D' differing by one record, and for any subset of outputs $S \subseteq \mathcal{R}$, we have:*

$$\Pr[\mathcal{M}(D) \in S] \leq e^{\epsilon} \times \Pr[\mathcal{M}(D') \in S] + \delta.$$

where ϵ is the privacy budget and δ is the failure probability.

When $\delta = 0$ we achieve a strictly stronger notion of ϵ-differential privacy. A common way to achieve differential privacy is to add some randomized noise to the output, where the noise is proportional to sensitivity of function $\mathcal{M}$:

Definition 2 (Sensitivity). *The sensitivity of function $\mathcal{M}$ is the maximum change in the output of $\mathcal{M}$ for all possible pairs of neighboring inputs D, D':*

$$\Delta = \max_{D,D'} \|\mathcal{M}(D) - \mathcal{M}(D')\|.$$

2.2 Differentially Private ERM

In this paper, we consider the regularized ERM. Formally, let $\mathcal{X}$ denote the input space and $\mathcal{Y}$ denote the output space, given a dataset $\mathcal{D} = \{\boldsymbol{z}_1, \boldsymbol{z}_2, \ldots, \boldsymbol{z}_n\}$ where $\boldsymbol{z}_i = (\boldsymbol{x}_i, y_i) \in \mathcal{X} \times \mathcal{Y}$, the objective of regularized ERM is to get a model $\hat{\boldsymbol{\theta}}$ from the following unconstrained optimization problem:

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} L(\boldsymbol{\theta}, \mathcal{D}) + \frac{\Lambda}{2n} ||\boldsymbol{\theta}||^2, \tag{1}$$

where $L(\boldsymbol{\theta}, \mathcal{D}) = \frac{1}{n} \sum_{i=1}^{n} \ell(\boldsymbol{\theta}, \boldsymbol{z}_i)$, $\ell(\boldsymbol{\theta}, \boldsymbol{z}_i)$ is the loss function for $\boldsymbol{z}_i$. The term of $\frac{\Lambda}{2n} ||\boldsymbol{\theta}||^2$ is the regularizer that prevents over-fitting.

There are 3 obvious candidates for where to add privacy-preserving noise for machine learning algorithms (see Figure 1). First, we could add noise to the objective function, which gives us the objective perturbation mechanism. Second, we could add noise to the gradients at each iteration, which gives us the gradient perturbation mechanism. Finally, we can add noise to the model after the training process, which gives us the output perturbation mechanism. In this paper, we focus our research on objective perturbation mechanism among the three.

The objective perturbation mechanism was first proposed in Chaudhuri et al. [2011] to make ERM algorithms achieving ϵ-differential privacy. Later, the range of problems to which objective perturbation applies has been extended in Kifer et al. [2012], including convex ERM problems, non-differentiable regularizers and tighter error bound.

Theorem 1 (Private Convex Optimization via Objective Perturbation in Kifer et al. [2012]). *Let $\mathcal{C}$ be a closed convex subset of $\mathbb{R}^p$, $\mathcal{D} = \{z_1, z_2, \ldots, z_n\}$ be a dataset, $\ell(\boldsymbol{\theta}, z_i)$ be convex loss function with continuous Hessian, $||\nabla \ell(\boldsymbol{\theta}, z_i)|| \le \zeta$ and $||\nabla^2 \ell(\boldsymbol{\theta}, z_i)|| \le \lambda$ (for all z and for all $\boldsymbol{\theta} \in \mathcal{C}$), $L(\boldsymbol{\theta}, \mathcal{D}) = \frac{1}{n} \sum_{i=1}^n \ell(\boldsymbol{\theta}, z_i)$, $\Delta_\epsilon = \frac{2\lambda}{\epsilon}$. Then*

$$\hat{\boldsymbol{\theta}}_\epsilon = \arg\min_{\boldsymbol{\theta}} L(\boldsymbol{\theta}, \mathcal{D}) + \frac{\Lambda}{2n}||\boldsymbol{\theta}||^2 + \frac{1}{n}\mathbf{b}_\epsilon^T \boldsymbol{\theta} + \frac{\Delta_\epsilon}{2n}||\boldsymbol{\theta}||^2 \tag{2}$$

is (ϵ, δ)-differentially private if $\mathbf{b}_\epsilon \sim \mathcal{N}\left(0, \frac{\zeta^2(8\log\frac{2}{\delta}+4\epsilon)}{\epsilon^2}\boldsymbol{I}\right)$.

Remark 1. *The reasons why we choose objective perturbation (not output perturbation or gradient perturbation) as the basis of our analysis are as follows: a) as one of the most important mechanisms for differentially private algorithms, objective perturbation has shown superior performance on general objectives Chaudhuri et al. [2011], Kifer et al. [2012], Iyengar et al. [2019]; b) for objective perturbation, our analysis of utility and ϵ can be directly established on the objective function, which would make the process intuitive and concise.*

3   Traditional Utility Guarantee Analysis in Private ERM

The exploration of the effect of ϵ on utility in private ERM started from the very beginning in Chaudhuri et al. [2011], where the differentially private ERM algorithms were first proposed. They found that differential privacy not only provides privacy guarantee, but also reduces the model utility. They measure the utility by the number of samples n required to achieve generalization error ξ on the empirical loss of a reference predictor $\boldsymbol{\theta}_0$, and bound n by using the analysis method of Shalev-Shwartz and Srebro [2008]. Their bounds provide the sample requirement of differentially private algorithms for achieving generalization error. That is, the algorithms will achieve higher utility with more training data, and as long as there is abundant training data, the utility of differentially private algorithms can be very close to non-private one. They also show that the error rate increases as the privacy requirements are made more stringent, which they call the "price of privacy". This is the first explicit reference to the tradeoff between privacy and utility of differentially private learning algorithms.

However, this form of expression is difficult to intuitively understand how ϵ effects the model utility, neither the utility loss nor the empirical loss can be acquired. To this end, Kifer et al. [2012] provided an improved, data-dependent utility analysis, which used empirical risk $J(\boldsymbol{\theta}^{priv}, \mathcal{D}) - J(\hat{\boldsymbol{\theta}}, \mathcal{D})$ to measure the utility, where $\hat{\boldsymbol{\theta}}$ denotes the optimal predictor, $\boldsymbol{\theta}^{priv}$ denotes the private predictor and $J(\boldsymbol{\theta}, \mathcal{D})$ denotes the empirical loss. By using the tail bounds for the noise distributions used in Theorem 1, they obtain the following empirical risk bound for their objective perturbed learning algorithm.

Theorem 2 (Empirical Risk Bound of Differentially Private Learning Algorithm in Kifer et al. [2012]). *Let $\hat{\boldsymbol{\theta}}$ be the true minimizer of $J(\boldsymbol{\theta}, \mathcal{D})$ over the closed convex set $\mathcal{C}$ and let $\hat{\boldsymbol{\theta}}_\epsilon$ be the output of (2) in Theorem 1. We have* $\mathbb{E}\left[J(\hat{\boldsymbol{\theta}}_\epsilon, \mathcal{D}) - J(\hat{\boldsymbol{\theta}}, \mathcal{D})\right] = O\left(\frac{\zeta||\hat{\boldsymbol{\theta}}||\sqrt{p\log(1/\delta)}}{\epsilon n}\right)$.

Theorem 2 clearly illustrates the upper bound of utility loss in the worst case. This method and its adaptions have been widely adopted up to the present literatures Bassily et al. [2014], Wang et al. [2017], Iyengar et al. [2019] for measuring the utility of private learning algorithms. Although these "utility theorems" can give an intuitive impression of the effect of ϵ on utility, their bounds are too conservative to find an appropriate value of ϵ (usually leading to a much larger value of ϵ). This makes them difficult to do something useful in practical applications.

In order to overcome this conservatism and give a more realistic analysis of the effect of ϵ on utility, we propose an "empirical" theorem approach.

4   Our Approach

In this section, we will introduce our approach for analyzing the effect of ϵ on the utility with gradual steps.

First of all, we formalize our final goal by asking the following question: how would the utility of the private model change if we modify the value of ϵ? To answer the question, we will begin by studying the change in model parameters due to modifying the value of ϵ. Formally, this change is $\hat{\boldsymbol{\theta}}_{\bar{\epsilon}} - \hat{\boldsymbol{\theta}}_{\hat{\epsilon}}$, where $\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}$ and $\hat{\boldsymbol{\theta}}_{\hat{\epsilon}}$ are the true minimizers of Equation

(2) that achieves $(\bar{\epsilon}, \delta)$- and $(\hat{\epsilon}, \delta)$-differential privacy on the same dataset $\mathcal{D}$, respectively, for a certain δ. Apparently, our analysis focuses on the influence of ϵ on the parameters $\hat{\boldsymbol{\theta}}_\epsilon$, which is $\frac{\partial \hat{\boldsymbol{\theta}}_\epsilon}{\partial \epsilon}$.

Theorem 3. *Let* $L(\boldsymbol{\theta}, \mathcal{D}), \Delta_\epsilon$ *and* $\hat{\boldsymbol{\theta}}_\epsilon$ *be defined as in Theorem 1. Let* $\boldsymbol{W}_\epsilon = \nabla^2 L(\hat{\boldsymbol{\theta}}_\epsilon, \mathcal{D}) + \frac{\Lambda}{n} + \frac{1}{n}\Delta_\epsilon, \mathbf{b}'_\epsilon = \frac{\partial \mathbf{b}_\epsilon}{\partial \epsilon}, \Delta'_\epsilon = \frac{\partial \Delta_\epsilon}{\partial \epsilon}$. *Then, we have*

$$\frac{\partial \hat{\boldsymbol{\theta}}_\epsilon}{\partial \epsilon} = -\frac{1}{n}\boldsymbol{W}_\epsilon^{-1}\left(\mathbf{b}'_\epsilon + \Delta'_\epsilon \hat{\boldsymbol{\theta}}_\epsilon\right). \tag{3}$$

Proof. Due to the fact that the derivative of the objective function (2) at $\hat{\boldsymbol{\theta}}_\epsilon$ is $\mathbf{0}$, we have that

$$\frac{\partial L(\hat{\boldsymbol{\theta}}_\epsilon, \mathcal{D})}{\partial \boldsymbol{\theta}} + \frac{1}{n}\left(\Lambda \hat{\boldsymbol{\theta}}_\epsilon + \mathbf{b}_\epsilon + \Delta_\epsilon \hat{\boldsymbol{\theta}}_\epsilon\right) = \mathbf{0}. \tag{4}$$

Take the derivative of (4) with respect to ϵ, we have:

$$\frac{\partial^2 L(\hat{\boldsymbol{\theta}}_\epsilon, \mathcal{D})}{\partial \boldsymbol{\theta}^2}\frac{\partial \hat{\boldsymbol{\theta}}_\epsilon}{\partial \epsilon} + \frac{1}{n}\left(\Lambda \frac{\partial \hat{\boldsymbol{\theta}}_\epsilon}{\partial \epsilon} + \mathbf{b}'_\epsilon + \frac{\partial(\Delta_\epsilon \hat{\boldsymbol{\theta}}_\epsilon)}{\partial \epsilon}\right) = \mathbf{0}.$$

Re-arranging the terms, we get

$$\left(\nabla^2 L(\hat{\boldsymbol{\theta}}_\epsilon, \mathcal{D}) + \frac{\Lambda}{n} + \frac{1}{n}\Delta_\epsilon\right)\frac{\partial \hat{\boldsymbol{\theta}}_\epsilon}{\partial \epsilon} = -\frac{1}{n}\left(\mathbf{b}'_\epsilon + \Delta'_\epsilon \hat{\boldsymbol{\theta}}_\epsilon\right).$$

Finally, we have

$$\frac{\partial \hat{\boldsymbol{\theta}}_\epsilon}{\partial \epsilon} = -\frac{1}{n}\boldsymbol{W}_\epsilon^{-1}\left(\mathbf{b}'_\epsilon + \Delta'_\epsilon \hat{\boldsymbol{\theta}}_\epsilon\right). \tag{5}$$

$\square$

To obtain $\frac{\partial \hat{\boldsymbol{\theta}}_\epsilon}{\partial \epsilon}$, we need to derive $\mathbf{b}_\epsilon$ over ϵ. However, since $\mathbf{b}_\epsilon$ doesn't have an explicit function expression, we can not compute $\frac{\partial \mathbf{b}_\epsilon}{\partial \epsilon}$ directly. Thus, we need to first form a function that outputs $\mathbf{b}_\epsilon$. To this end, we make use of the reparameterization trick, which is commonly applied to training variational autoencoders with continuous latent variables using backpropagation Kingma and Welling [2014], Jimenez Rezende et al. [2014]. For a given normal distribution $z \sim \mathcal{N}(\mu, \sigma^2)$, it can be re-written as $z = \mu + \sigma \cdot \mathcal{N}(0, 1)$ by reparameterization, making it trivial to compute $\frac{\partial z}{\partial \sigma}$ and $\frac{\partial z}{\partial \mu}$.

Now we will show how to use the reparameterization trick to compute $\frac{\partial \mathbf{b}_\epsilon}{\partial \epsilon}$. Let $\mathbf{b}_\epsilon = (\mathbf{b}_1, \ldots, \mathbf{b}_d)$, then $\mathbf{b}'_\epsilon = (\mathbf{b}'_1, \ldots, \mathbf{b}'_d)$. Since $\mathbf{b}_\epsilon \sim \mathcal{N}\left(0, \frac{\zeta^2(8\log\frac{2}{\delta}+4\epsilon)}{\epsilon^2}\boldsymbol{I}\right)$, we have that for each $i \in [d]$, $\mathbf{b}_i$ is independently sampled from $\mathcal{N}\left(0, \frac{\zeta^2(8\log\frac{2}{\delta}+4\epsilon)}{\epsilon^2}\right)$. By using reparameterization trick, we obtain $\mathbf{b}_i = \left(\frac{\zeta\sqrt{8\log\frac{2}{\delta}+4\epsilon}}{\epsilon}\right) \cdot \mathcal{N}(0, 1)$. Note that $\mathcal{N}(0, 1)$ can be seen as a constant term for ϵ, we have that $\mathbf{b}'_i = \left(\frac{\zeta(4\log\frac{2}{\delta}+\epsilon)}{\epsilon^2\sqrt{2\log\frac{2}{\delta}+\epsilon}}\right) \cdot \mathcal{N}(0, 1)$. Thus, we have:

$$\mathbf{b}'_\epsilon = \left(\frac{\zeta\left(4\log\frac{2}{\delta} + \epsilon\right)}{\epsilon^2\sqrt{2\log\frac{2}{\delta} + \epsilon}}\right) \cdot \mathcal{N}(0, \boldsymbol{I}). \tag{6}$$

The assumptions in Theorem 3 can be satisfied on many commonly used loss functions, such as logistic regression, Huber SVM and quadratic loss, and regularizers, such as L_2-norm regularizer. However, this does exclude some classes of loss functions (e.g. hinge loss). In Section 5.1, we will talk about this cases and show that some of these assumptions can be weakened in certain cases.

Once $\frac{\partial \hat{\boldsymbol{\theta}}_\epsilon}{\partial \epsilon}$ is calculated, we can further deduce the effect of ϵ on the utility. In particular, let $F(\boldsymbol{\theta}, \mathcal{D})$ be the utility of $\boldsymbol{\theta}$ on dataset $\mathcal{D}$, then, by applying the chain rule, we have the following expression:

$$\frac{\partial F(\hat{\boldsymbol{\theta}}_\epsilon, \mathcal{D})}{\partial \epsilon} = \left(\nabla_{\boldsymbol{\theta}} F(\hat{\boldsymbol{\theta}}_\epsilon, \mathcal{D})\right)^T \frac{\partial \hat{\boldsymbol{\theta}}_\epsilon}{\partial \epsilon} \tag{7}$$

Now we can answer the question posed at the beginning of this section by utilizing the Taylor expansion of $F(\bar{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D}) - F(\hat{\boldsymbol{\theta}}_{\hat{\epsilon}}, \mathcal{D})$ at $\bar{\epsilon}$:

$$F(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D}) - F(\hat{\boldsymbol{\theta}}_{\hat{\epsilon}}, \mathcal{D}) \approx \frac{\partial F(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D})}{\partial \epsilon}(\bar{\epsilon} - \hat{\epsilon}). \tag{8}$$

Intuitively, equation (8) formulizes the relationship between the change in ϵ and the change in utility. Furthermore, once we have obtained $F(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D})$ with a certain $\bar{\epsilon}$ (through one round training), (8) represents the relationship between $\hat{\epsilon}$ and its utility, which could be used for choosing ϵ according to expected utility. This application would be further discussed in Section 6. Generally, $\bar{\epsilon}$ in (8) is called measuring point, and $\hat{\epsilon}$ in (8) is called target point.

In essence, both $F(\hat{\boldsymbol{\theta}}_{\epsilon}, \mathcal{D})$ in (7) and $L(\hat{\boldsymbol{\theta}}_{\epsilon}, \mathcal{D})$ in (3) are used to measure the error between predictions and true labels, which sometimes may even be the same in applications. Thus, for the simplicity of exposition, we assume that all the assumptions in $L(\hat{\boldsymbol{\theta}}_{\epsilon}, \mathcal{D})$ are equally valid in $F(\hat{\boldsymbol{\theta}}_{\epsilon}, \mathcal{D})$.

5 Practical Considerations

In this section, we are going to talk about two practical considerations of our approximation approach.

5.1 Assumption Violation

Theorem 3 relies on several assumptions that may be violated in practice. The first would be non-convex or non-convergent objective functions. In this case, the obtained parameters $\tilde{\boldsymbol{\theta}}_{\epsilon}$ may not be the global minimum, thus Equation (4) may not hold. In such cases, we can form a convex quadratic approximation of the loss around $\tilde{\boldsymbol{\theta}}_{\epsilon}$, i.e.,

$$L(\boldsymbol{\theta}, \mathcal{D}) \approx L(\tilde{\boldsymbol{\theta}}_{\epsilon}, \mathcal{D}) + \nabla_{\boldsymbol{\theta}} L(\tilde{\boldsymbol{\theta}}_{\epsilon}, \mathcal{D})(\tilde{\boldsymbol{\theta}}_{\epsilon} - \boldsymbol{\theta}) + (\tilde{\boldsymbol{\theta}}_{\epsilon} - \boldsymbol{\theta})^T \nabla_{\boldsymbol{\theta}}^2 L(\tilde{\boldsymbol{\theta}}_{\epsilon}, \mathcal{D})(\tilde{\boldsymbol{\theta}}_{\epsilon} - \boldsymbol{\theta}),$$

which corresponds to adding L_2 regularization on $\boldsymbol{\theta}$.

The second violation would be non-differentiable losses, in which neither $\nabla_{\boldsymbol{\theta}} L(\hat{\boldsymbol{\theta}}_{\epsilon}, \mathcal{D})$ nor $\nabla_{\boldsymbol{\theta}}^2 L(\hat{\boldsymbol{\theta}}_{\epsilon}, \mathcal{D})$ does exist. In these cases of violation, we can train non-differentiable models and swap out non-differentiable components for smoothed versions (as suggested in Koh and Liang [2017]). For example, we can approximated $Hinge(s) = max(0, 1 - s)$ with following form $SmoothHinge(s, t) = t \log(1 + e^{\frac{1-s}{t}})$, which approaches the hinge loss at $t \to 0$.

Remark 2. *The above analysis shows that our approximate approach can be extended to the non-convex, non-convergent and non-differentiable situations. However, in this paper, we mainly focus on verifying the effectiveness of our approach. Therefore, we only consider the normal case in our evaluations. We leave the important analysis of our approximate approach for non-convex, non-convergent and non-differentiable cases to our future work.*

5.2 Error Analysis

To analyze the error bound of our approach, we need to complete the Taylor expression of Equation (8) as follows:

$$F(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D}) - F(\hat{\boldsymbol{\theta}}_{\hat{\epsilon}}, \mathcal{D}) \approx \frac{\partial F(\hat{\boldsymbol{\theta}}_{\epsilon}, \mathcal{D})}{\partial \epsilon}(\bar{\epsilon} - \hat{\epsilon}) + r(\bar{\epsilon} - \hat{\epsilon}), \tag{9}$$

where $r(\bar{\epsilon} - \hat{\epsilon}) = \frac{\partial^2 F(\hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}, \mathcal{D})}{\partial \epsilon^2}(\bar{\epsilon} - \hat{\epsilon})^2$ is the Taylor remainder, $\tilde{\epsilon} \in [\bar{\epsilon}, \hat{\epsilon}]$. Obviously, our error mainly comes from omitting the Taylor remainder. Thus, to derive the error bound of Equation (8), we require the Taylor remainder to be bounded.

Theorem 4. *Assume that $\ell(\boldsymbol{\theta}, \mathcal{D})$ is third-differentiable and convex in $\boldsymbol{\theta}$ with continuous Hessian, $||\nabla \ell(\boldsymbol{\theta}, \boldsymbol{z}_i)|| \leq c$, $||\nabla^2 \ell(\boldsymbol{\theta}, \boldsymbol{z}_i)|| \leq \lambda$, $\Delta_{\epsilon} = \frac{2\lambda}{\epsilon}$, $\mathbf{b}_{\epsilon} \sim \mathcal{N}\left(0, \frac{c^2(8 \log \frac{2}{\delta} + 4\epsilon)}{\epsilon^2}\boldsymbol{I}\right)$. If there exists a constant s for $\forall \boldsymbol{z} \in \mathcal{D}, \boldsymbol{\theta} \in \mathcal{C}, ||\nabla_{\boldsymbol{\theta}}^3 \ell(\boldsymbol{\theta}, \mathcal{D})|| \leq s$. Then, for any $\bar{\epsilon}, \hat{\epsilon}$ we have*

$$|F(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D}) - F(\hat{\boldsymbol{\theta}}_{\hat{\epsilon}}, \mathcal{D})| = \mathcal{O}\left(\frac{(\bar{\epsilon} - \hat{\epsilon})^2}{n\tilde{\epsilon}^3}\right) \tag{10}$$

where $\tilde{\epsilon} \in [\bar{\epsilon}, \hat{\epsilon}]$, n is the number of training samples. $\mathcal{O}$ omitted the terms p, λ, δ, c, s.

Proof. Let's try expanding $\frac{\partial^2 F(\hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}, \mathcal{D})}{\partial \epsilon^2}$ for the first time:

$$
\begin{aligned}
\frac{\partial^2 F(\hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}, \mathcal{D})}{\partial \epsilon^2} &= \frac{\partial \left(\frac{\partial F(\hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}, \mathcal{D})}{\partial \epsilon} \right)}{\partial \epsilon} \\
&= \frac{\partial \left(\left(\frac{\partial F(\hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}, \mathcal{D})}{\partial \boldsymbol{\theta}} \right)^T \frac{\partial \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon} \right)}{\partial \epsilon} \\
&= \left(\frac{\partial^2 F(\hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}, \mathcal{D})}{\partial \boldsymbol{\theta} \partial \epsilon} \right)^T \frac{\partial \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon} + \left(\frac{\partial F(\hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}, \mathcal{D})}{\partial \boldsymbol{\theta}} \right)^T \frac{\partial^2 \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon^2} \\
&= \left(\frac{\partial \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon} \right)^T \left(\nabla^2 F(\hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}, \mathcal{D}) \right) \frac{\partial \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon} + \left(\nabla F(\hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}, \mathcal{D}) \right)^T \frac{\partial^2 \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon^2}.
\end{aligned}
\tag{11}
$$

To go on, we need to expand $\frac{\partial^2 \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon^2}$, which is as follows:

$$
\begin{aligned}
\frac{\partial^2 \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon^2} &= \frac{\partial \left(\frac{\partial \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon} \right)}{\partial \epsilon} \\
&= \frac{\partial \left(-\frac{1}{n} \boldsymbol{W}_{\tilde{\epsilon}}^{-1} \left(\mathbf{b}'_{\tilde{\epsilon}} + \Delta'_{\tilde{\epsilon}} \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}} \right) \right)}{\partial \epsilon} \\
&= -\frac{1}{n} \frac{\partial \left(\boldsymbol{W}_{\tilde{\epsilon}}^{-1} \right)}{\partial \epsilon} \left(\mathbf{b}'_{\tilde{\epsilon}} + \Delta'_{\tilde{\epsilon}} \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}} \right) - \frac{1}{n} \boldsymbol{W}_{\tilde{\epsilon}}^{-1} \frac{\partial \left(\mathbf{b}'_{\tilde{\epsilon}} + \Delta'_{\tilde{\epsilon}} \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}} \right)}{\partial \epsilon} \\
&= \boldsymbol{W}_{\tilde{\epsilon}}^{-1} \frac{\partial \boldsymbol{W}_{\tilde{\epsilon}}}{\partial \epsilon} \left(-\frac{1}{n} \right) \boldsymbol{W}_{\tilde{\epsilon}}^{-1} \left(\mathbf{b}'_{\tilde{\epsilon}} + \Delta'_{\tilde{\epsilon}} \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}} \right) - \frac{1}{n} \boldsymbol{W}_{\tilde{\epsilon}}^{-1} \left(\frac{\partial (\mathbf{b}'_{\tilde{\epsilon}})}{\partial \epsilon} + \frac{\partial (\Delta'_{\tilde{\epsilon}})}{\partial \epsilon} \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}} + \Delta'_{\tilde{\epsilon}} \frac{\partial \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon} \right) \\
&= \boldsymbol{W}_{\tilde{\epsilon}}^{-1} \left(\frac{\partial \left(\nabla^2 L(\hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}, \mathcal{D}) \right)}{\partial \epsilon} + \frac{1}{n} \Delta'_{\tilde{\epsilon}} \boldsymbol{I} \right) \frac{\partial \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon} - \frac{1}{n} \boldsymbol{W}_{\tilde{\epsilon}}^{-1} \left(\mathbf{b}''_{\tilde{\epsilon}} + \Delta''_{\tilde{\epsilon}} \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}} + \Delta'_{\tilde{\epsilon}} \frac{\partial \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon} \right) \\
&= \boldsymbol{W}_{\tilde{\epsilon}}^{-1} \left(\nabla^3 L(\hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}, \mathcal{D}) \right) \left(\frac{\partial \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon} \right) \left(\frac{\partial \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}}}{\partial \epsilon} \right) - \frac{1}{n} \boldsymbol{W}_{\tilde{\epsilon}}^{-1} \left(\mathbf{b}''_{\tilde{\epsilon}} + \Delta''_{\tilde{\epsilon}} \hat{\boldsymbol{\theta}}_{\tilde{\epsilon}} \right),
\end{aligned}
\tag{12}
$$

where $\mathbf{b}''_{\tilde{\epsilon}} = \frac{\partial (\mathbf{b}'_{\tilde{\epsilon}})}{\partial \epsilon}, \Delta''_{\tilde{\epsilon}} = \frac{\partial (\Delta'_{\tilde{\epsilon}})}{\partial \epsilon}$.

Now, let's bound the term $||\boldsymbol{W}_{\tilde{\epsilon}}^{-1}||$ first. By the assumptions in Theorem 3, we have

$$
\begin{aligned}
||\boldsymbol{W}_{\tilde{\epsilon}}|| &= ||\nabla^2 L(\boldsymbol{\theta}, \mathcal{D})|| + \frac{\Lambda}{n} + \frac{1}{n}||\Delta_\epsilon|| \\
&= \frac{1}{n} \sum_{i=1}^{n} ||\nabla^2 \ell(\boldsymbol{\theta}, \boldsymbol{z}_i)|| + \frac{\Lambda}{n} + \frac{2\lambda}{n\epsilon} \\
&= \mathcal{O} \left(\lambda + \frac{\Lambda}{n} + \frac{2\lambda}{n\epsilon} \right).
\end{aligned}
\tag{13}
$$

Thus, $||\boldsymbol{W}_{\bar{\epsilon}}^{-1}|| = \mathcal{O}\left(\frac{1}{\lambda + \frac{\Delta}{n} + \frac{2\lambda}{n\epsilon}}\right)$. Since it is usally that $\epsilon \gg \frac{1}{n}$, we obtain the bound of $\boldsymbol{W}_{\bar{\epsilon}}^{-1}$ to be $\mathcal{O}\left(\frac{1}{\lambda}\right)$.

As $\Delta_{\epsilon} = \frac{2\lambda}{\epsilon}$, we have that $\Delta'_{\epsilon} = \mathcal{O}\left(\frac{\lambda}{\epsilon^2}\right)$ and $\Delta''_{\epsilon} = \mathcal{O}\left(\frac{\lambda}{\epsilon^3}\right)$.

Then we bound $\mathbf{b}'_{\bar{\epsilon}}$ and $\mathbf{b}''_{\bar{\epsilon}}$. We find $\log\frac{2}{\delta} + \epsilon$ and $\log\frac{2}{\delta}$ are of the same order of magnitude as large ϵ values are rarely used. Thus we have that $\mathbf{b}'_{\bar{\epsilon}} = \mathcal{O}\left(\frac{\sqrt{\log(\frac{1}{\delta})}}{\epsilon^2}\right)$ acoording to Equation (6). Similarly, we have that $\mathbf{b}''_{\bar{\epsilon}} = \mathcal{O}\left(\frac{\sqrt{\log(\frac{1}{\delta})}}{\epsilon^3}\right)$.

Note that $\frac{\partial\hat{\boldsymbol{\theta}}_{\epsilon}}{\partial\epsilon}$ can be easily calculated by Equation (5), $||\boldsymbol{\theta}||$ is usually assumed to be at most $\sqrt{p}$ via normalization (p is the dimensionality of data, which is not considered in our paper), $||\nabla L(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D})|| \le c$, $||\nabla^2 L(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D})|| \le \lambda$ and $||\nabla^3 L(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D})|| \le s$, we can obtain our error bound of $r(\bar{\epsilon} - \hat{\epsilon}) = \frac{\partial^2 L(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D})}{\partial\epsilon^2}(\bar{\epsilon} - \hat{\epsilon})^2$ as follows:

$$\mathcal{O}\left(\frac{(\bar{\epsilon} - \hat{\epsilon})^2}{n\tilde{\epsilon}^3}\right),$$

where $\mathcal{O}$ omitted the terms p, λ, δ, c, s. □

To bound the Taylor remainder, we assume the loss function has bounded third derivative. For most commonly used twice-differentiable loss functions, this assumption is easy to be satisfied. For example, the third derivative of quadratic loss is zero, any order derivative of logistic regression is bounded, etc..

Theorem 4 shows that our error bound is proportional to the difference between $\bar{\epsilon}$ and $\hat{\epsilon}$, and inversely proportional to sample number n. That is, the more samples we used for training and the smaller the change in ϵ, the higher estimation accuracy we could obtain. Moreover, the magnitude of $\tilde{\epsilon}$ (i.e. the magnitude of $\bar{\epsilon}$ and $\hat{\epsilon}$) is also an important factor affecting estimation accuracy. Considering the general cases, where $\bar{\epsilon}$ and $\hat{\epsilon}$ are of the same magnitude, our bound would approach to $\mathcal{O}\left(\frac{1}{n\tilde{\epsilon}}\right)$, which is fairly small when $\tilde{\epsilon}$ is of general magnitude (i.e. $10^{-2} \sim 10^1$). However, if $\bar{\epsilon}$ and $\hat{\epsilon}$ are of different magnitudes, especially when one is of small magnitude, the error may be relatively large. For example, when $\bar{\epsilon} = 0.01, \hat{\epsilon} = 10$, we have $(\bar{\epsilon} - \hat{\epsilon})^2 \approx 100$ and $\tilde{\epsilon} \in [0.01, 10]$, then the error would be between $\mathcal{O}(\frac{10^{-1}}{n})$ and $\mathcal{O}(\frac{10^8}{n})$. Thus, to obtain small error, we should ensure that $\bar{\epsilon}$ and $\hat{\epsilon}$ are of the same magnitude when using our approximation approach.

6   Practical Use Case of Our Approach

By revealing insights about how ϵ affects the model utility in differentially private ERM, our approach would be very useful in the practical application. In this section, we will show a typical use case about how our approach guides users to choose an appropriate ϵ for private learning in practice.

The question of how to choose ϵ is a typical application scenario, which has existed since differential privacy was first proposed Lee and Clifton [2011], Hsu et al. [2014]. Traditional approaches to differential privacy assume a fixed privacy requirement ϵ for the user and attempt to maximize the utility of model to the privacy constraint Naldi and Dacquisto [2015], Hsu et al. [2014], Kohli and Laskowski [2018]. However, existing approaches share a common precondition that the user is capable to perceive and judge the effects of a privacy parameter. Considering that the privacy is hard to be quantitative variable, and the definition of differential privacy offers little insight into how this should be done, it may be a little harsh to require users to have such ability.

Different from the existing approaches, we provide a more efficient and accurate approach to addressing this issue, based on the conclusions of the previous analysis. Conceptually, by re-arranging the terms in (8), we have:

$$\hat{\epsilon} \approx \frac{F(\hat{\boldsymbol{\theta}}_{\hat{\epsilon}}, \mathcal{D}) - F(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D})}{\frac{\partial F(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D})}{\partial\epsilon}} + \bar{\epsilon}. \tag{14}$$

Taking $F(\hat{\boldsymbol{\theta}}_{\hat{\epsilon}}, \mathcal{D})$ as the expected utility we want to satisfy, then $\hat{\epsilon}$ would be the approximation of the ϵ we are looking for (i.e. the minimum ϵ achieves $F(\hat{\boldsymbol{\theta}}_{\hat{\epsilon}}, \mathcal{D})$). Thus, once we obtain $F(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D})$ and $\frac{\partial F(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D})}{\partial\epsilon}$ by training a model $\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}$ with $\bar{\epsilon}$, $\hat{\epsilon}$ can be directly calculated without repeated trainings. Algorithm 1 gives a brief outline of this process.

Our approach mainly has two advantages for choosing ϵ. Compared with utility theorem, our approach can find nearly optimal ϵ that is very close to the true minimum value consistent with utility requirement as we analyzed in Section 5.2. And, compared with empirical search, our approach is much more efficient as we only need to train the algorithm on the full dataset once. This feature will play a great advantage in case large-scale learning tasks.

Algorithm 1 Choosing ϵ with our approximation approach

Input: Dataset $\mathcal{D} = \{\boldsymbol{z}_1, \boldsymbol{z}_2, \ldots, \boldsymbol{z}_n\}$, loss function $\ell(\boldsymbol{\theta}, \boldsymbol{z}_i)$, measuring privacy parameter $\bar{\epsilon}$, failure probability δ, loss
    function $\ell(\boldsymbol{\theta}, \boldsymbol{z}_i)$, expected utility $\hat{F}$.
 1: Train a private model $\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}$ on $\mathcal{D}$ by Theorem 1.
 2: Compute $\frac{\partial F(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D})}{\partial \epsilon}$ by Theorem 3
Output: $\hat{\epsilon} = \frac{\hat{F} - F(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D})}{\frac{\partial F(\hat{\boldsymbol{\theta}}_{\bar{\epsilon}}, \mathcal{D})}{\partial \epsilon}} + \bar{\epsilon}$

Table 1: Datasets used in our evaluations

Dataset	#Samples	#Dim.	#Classes
Adult	$45,220$	104	2
KDDCup99	$70,000$	114	2
Gisette	$10,000$	$5,000$	2

7 Evaluation

7.1 Evaluation Setup

In this section, we will empirically analyze the performance of our approach. We implement the objective perturbation mechanism of Kifer et al. [2012] based on the open source released by Iyengar et al. [2019]. Our evaluation considers the loss functions for two commonly used models: Logistic Regression (LR) and Huber SVM (SVM), which are defined as follows:

$$\ell_{LR}(\boldsymbol{z}) = \log(1 + e^{-\boldsymbol{z}}),$$

$$\ell_{SVM}(\boldsymbol{z}) = \begin{cases} 0, & \text{if } z > 1 + h, \\ \frac{1}{4h}(1 + h - z)^2, & \text{if } |1 - z| \leq h, \\ 1 - z, & \text{if } z < 1 - h, \end{cases} \tag{15}$$

both of which are twice-differentiable. The term h in SVM is set to be 0.1 by default. The hyperparameter Λ of regularizer is tuned by training a non-private model on each dataset and set to be 10^{-2} for best performance. We use stochastic gradient descent algorithm to minimize Equation (2), of which the iterations and learning rate are set to be 100 and 0.01. The ϵ values we choose for training private models are between 0 and 10, which is very commonly used in related researches.

The Adult Dua and Graff [2017], Kddcup99 Stamper et al. [2010] and Gisette Guyon et al. [2004] datasets are used in our experiments. The details of datasets is shown in Table 1. Due to the random noise addition, all experiments are repeated ten times and the average results are reported. The upper bound of $||\nabla \ell(\boldsymbol{\theta}, \boldsymbol{z}_i)||$ and eigenvalues of $\nabla^2 \ell(\boldsymbol{\theta}, \boldsymbol{z}_i)$ is set to be $2\sqrt{p}$ and p due to Kifer et al. [2012], where p denotes the dimensionality of dataset.

7.2 Experiment Design

We totally set up three experiments. The first experiment is to verify the effectiveness of our approach. For each dataset, we randomly select $10,000$ samples and use them to train private models by the original objective perturbation mechanism according to each ϵ value, and use part of them as the measuring points to estimate empirical losses of other ϵ values according to Equation (8), then compare estimation results with the real results. Due to our expectation, the gap between real results and calculated results should be very small.

The second experiment is to analyze the impact of the selection of measuring points on the estimation results. We use each ϵ value as measuring point to estimate the empirical losses of all other ϵ values, and calculate the average error of estimated results. We hope that some rules can be found, which can be used as the guidance for selecting measuring point when using our approach.

The third experiment is to evaluate the effect of sample number on estimation error. Due to the limited sample number of Gisette, we only use Adult and Kddcup99 in this experiment. The number of samples is set from 1000 to 20000 for each dataset with fixed measuring points. According to our analysis in Section 5.2, the estimation error should decrease with the increase of sample number.

(a) Adult-LR

(b) Adult-SVM

(c) Kddcup99-LR

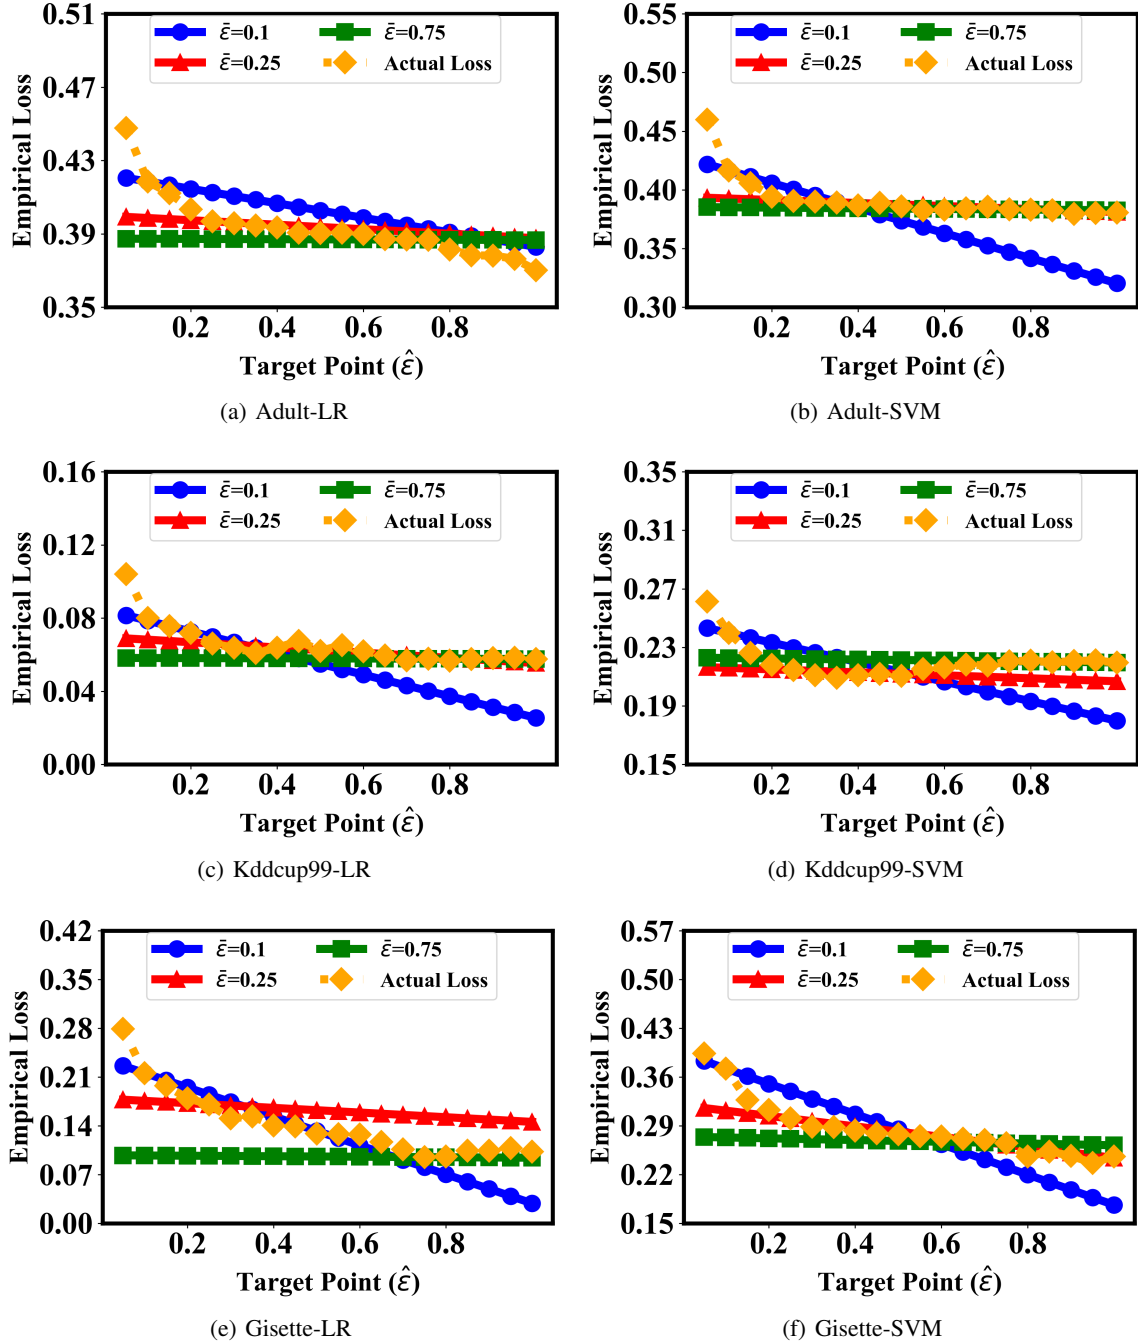(d) Kddcup99-SVM

(e) Gisette-LR

(f) Gisette-SVM

Figure 2: Performance of our approximation approach on datasets of Adult and Kddcup99 with logistic regression (LR) loss and Huber SVM (SVM) loss ($0 < \epsilon \leq 1.0$).
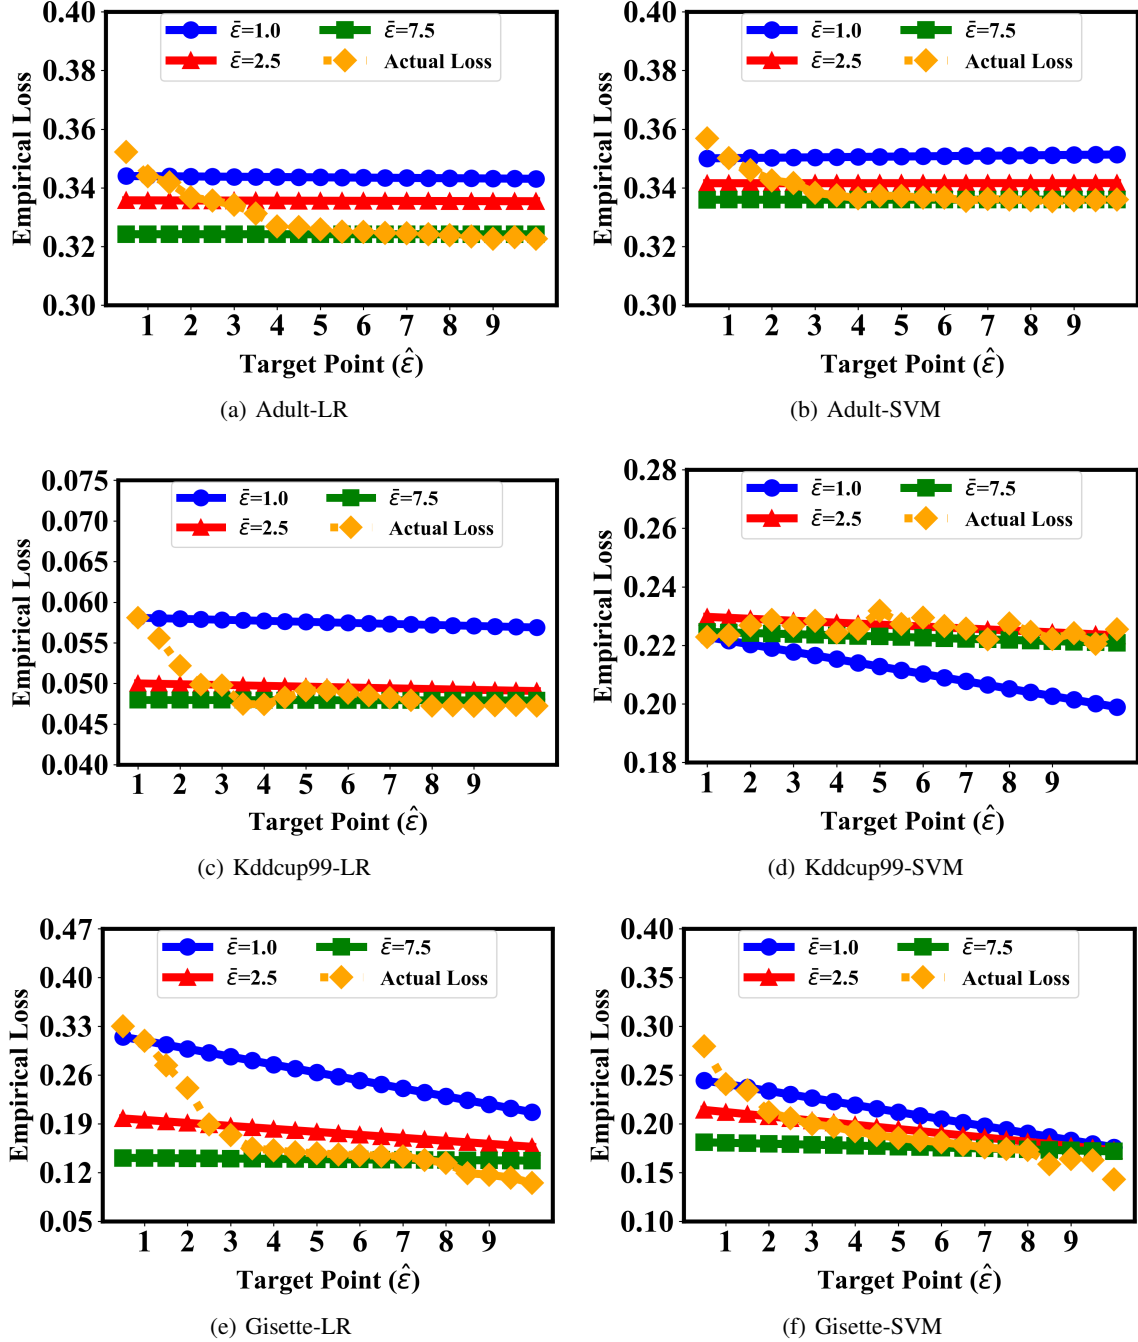
Figure 3: Performance of our approximation approach on datasets of Adult and Kddcup99 with logistic regression (LR) loss and Huber SVM (SVM) loss ($1.0 \leq \epsilon \leq 10$).

7.3 Evaluation Results and Analysis

To make the results of the first experiment more clear, we divide ϵ into two domains due to the magnitude, which are $0 < \epsilon \leq 1.0$ and $1.0 \leq \epsilon \leq 10$. For $0 < \epsilon \leq 1.0$, we choose $\bar{\epsilon} = 0.1, \bar{\epsilon} = 0.25, \bar{\epsilon} = 0.75$ as the measuring points to estimate the empirical loss at points varying from $\hat{\epsilon} = 0.05$ to $\hat{\epsilon} = 1.0$. For $1.0 \leq \epsilon \leq 10$, we choose $\bar{\epsilon} = 1.0, \bar{\epsilon} = 2.5, \bar{\epsilon} = 7.5$ as the measuring points to estimate the empirical loss at points varying from $\hat{\epsilon} = 1.0$ to $\hat{\epsilon} = 10$.

The experimental results of $0 < \epsilon \leq 1.0$ are shown in Figure 2, and the experimental results of $1.0 \leq \epsilon \leq 10$ are shown in Figure 3. For both figures, the orange line representing the real empirical loss of each ϵ is called the real line, and the other lines representing the estimate results of empirical loss are called estimated lines. Each estimated line is drawn by the estimation results of the empirical loss of each $\hat{\epsilon}$ (target points). Obviously, from both Figure 2 and Figure 3, it can be seen that the estimated lines basically fit the real line on each dataset (i.e. they are very close to each other). The results strongly confirm the effectiveness of our approach. Relatively speaking, the closer target point to measuring point, the higher estimation accuracy. This is basically consistent with our error analysis, that our error bound is proportional to the difference between $\bar{\epsilon}$ and $\hat{\epsilon}$. In addition, it can be seen that estimated lines are straight with different slope. The reason for this is that after training at the measuring point $\bar{\epsilon}$, the coefficients in (3) are determined, which makes a linear relationship between the value of target $\hat{\epsilon}$ and its estimated empirical loss. Compared with Figure 2 and Figure 3, we can also see that when measuring and estimating on larger ϵ, the error will be further reduced (note that the scale span of ordinate in Figure 3 is less than that in Figure 2), which is in good agreement with the error analysis in Section 5.2.

For the second experiment, we focus on the analysis of $0 < \epsilon \leq 1.0$. Since there is quite a bit of experimental results, we divide the large horizontal table into three small horizontal tables (which are Table 2, Table 3) to ensure that the results are clearly displayed in the page. For each dataset (each line), the result in bold represents the measuring $\bar{\epsilon}$ with minimum error. We can see that the minimum error is concentrated in the middle, and the maximum error is either at the smallest $\bar{\epsilon}$ (i.e. $\bar{\epsilon} = 0.05$) or at the largest $\bar{\epsilon}$ (i.e. $\bar{\epsilon} = 1.00$). The reason for this is that the average gap between intermediate value and other values is smaller, so the average error of them tend to be smaller. In addition, the smallest $\bar{\epsilon}$ tends to have larger average error. This is because that $\tilde{\epsilon}$ in Equation (10) would be small when measuring $\bar{\epsilon}$ is small, which would lead to high estimated error. Thus, in application, it would be better to select the measuring point near the mean value of target points so as to reduce the estimation error.

For the third experiment, we fix the value of $\bar{\epsilon}$ for each dataset to be the one that achieves minimum error and vary the sample number from 1000 to 20000 in the second experiment. The results are shown in Figure 4. Since the objects we compare are all estimation results, we directly use estimation error (instead of empirical loss we used in the first experiment) as the ordinate of each figure. Obviously, the estimation error decreases with the increase of sample number. This is consistent with our analysis in Section 5.2 as the error bound of our approach is inversely proportional to sample number n due to Equation (10). One possible explanation for this is that more training samples can improve the accuracy of the trained model and reduce the empirical loss at measuring $\bar{\epsilon}$. From Figure 4, we can see our error analysis is confirmed again, that the closer to measuring $\bar{\epsilon}$, the lower the estimation error of target $\hat{\epsilon}$, and under the same distance, the estimation error of larger target $\hat{\epsilon}$ is less than that of smaller target $\hat{\epsilon}$.

8 Related Works

To overcome the conservatism of "utility theorems", Ligett et al. [2017] proposed an empirical approximation method for utility analysis. In their method, they formalize the utility analysis as finding the minimum ϵ that meets a sufficient level of accuracy. They start from a very private ϵ, then gradually subtract the noise until the expected accuracy is achieved. The advantage of this method is that it focuses on the reality, not the theoretical worst-case. Therefore, their analysis is much closer to the real value than "utility theorems". However, an obvious drawback of this method is that it requires so many attempts of ϵ (noise subtraction), that would take a lot of computation and training time. This is often difficult to be accepted in practice, especially in large-scale learning tasks. Compared with Ligett et al. [2017], our approach can effectively overcome the shortcomings of the "utility theorems" and the method in Ligett et al. [2017]. On the one hand, our approach is a data-dependent theory, thus the analysis tends to be consistent with the actual situation. On the other hand, instead of repeated training to approach the target ϵ step by step, our approach only needs one-round training to obtain the target ϵ directly, which makes it much more efficient than Ligett et al. [2017].

The accuracy-first differential privacy has also been widely studied in differentially private query answering systems. For example, GUPT Mohan et al. [2012] offers a tool based on the sample-and-aggregate framework for differential privacy, which allows analysts to specify the target accuracy of the output, and compute privacy from it—or vice versa. PSI Gaboardi et al. [2016] offers to the data analyst an interface for selecting either the level of accuracy that she wants to reach, or the level of privacy she wants to impose. APEx Ge et al. [2019] allows data analysts to pose

Table 2: Average error of estimated loss for measuring point from $\bar{\epsilon} = 0.05$ to $\bar{\epsilon} = 0.50$ (part-1).

| $\bar{\epsilon}$	0.05	0.10	0.15	0.20	0.25	0.30	0.35	0.40	0.45	0.50
Adlut-LR	0.01982	0.00823	0.00908	0.00408	**0.00003**	0.00033	0.00004	0.00037	0.00293	0.00309
Adlut-SVM	0.08188	0.02052	0.00327	0.00315	0.00445	0.00367	0.00385	0.00558	**0.00298**	0.00585
Kddcup-LR	0.05884	0.02393	0.00930	0.00721	0.00763	0.00845	0.00960	0.00677	**0.00283**	0.00778
Kddcup-SVM	0.04735	0.00866	0.00798	0.00793	0.00851	0.01052	0.01143	0.00954	0.00848	0.00967
Gisette-LR	0.18427	0.01482	0.01181	0.02125	0.01931	**0.00302**	0.00810	0.00372	0.00394	0.01361
Gisette-SVM	0.16245	0.00571	0.02480	0.00822	0.00567	0.01023	**0.00550**	0.00712	0.00823	0.00823
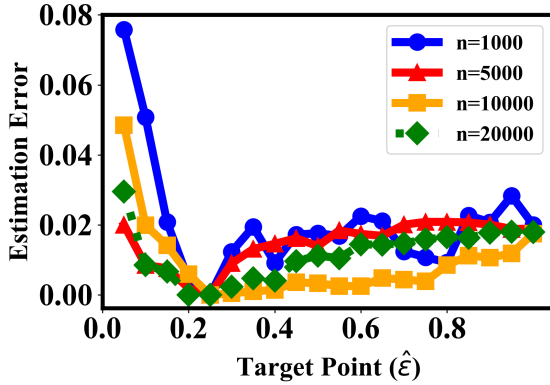
(a) Adult-LR
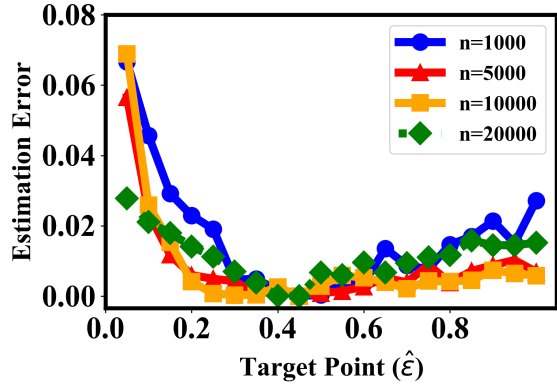
(b) Adult-SVM
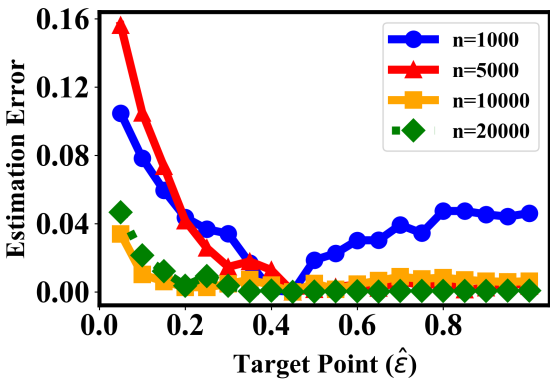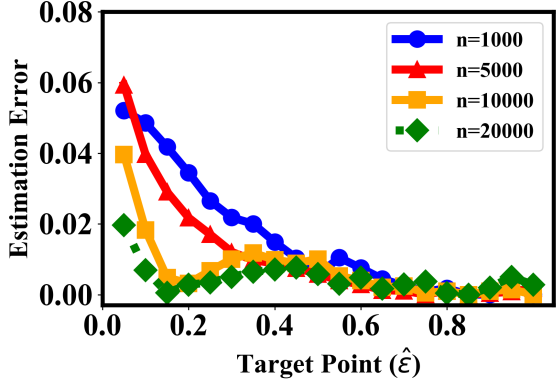
(c) Kddcup99-LR

(d) Kddcup99-SVM

Figure 4: Affect of sample number on estimation error. The measuring point for each evaluation, from left to right, is set to be $\bar{\epsilon} = 0.25, \bar{\epsilon} = 0.45, \bar{\epsilon} = 0.45, \bar{\epsilon} = 0.85$, respectively, which achieves the minimum error in the second experiment.

Table 3: Average error of estimated loss for measuring point from $\bar{\epsilon} = 0.55$ to $\bar{\epsilon} = 1$ (part-2).

$\bar{\epsilon}$	0.55	0.60	0.65	0.70	0.75	0.80	0.85	0.90	0.95	1.00
Adlut-LR	0.00339	0.00623	0.00643	0.00649	0.01169	0.01480	0.00283	0.01485	0.01675	0.02276
Adlut-SVM	0.00833	0.00711	0.00546	0.00770	0.00780	0.00840	0.00890	0.01142	0.01053	0.01020
Kddcup-LR	0.01186	0.01288	0.01200	0.01149	0.01147	0.01213	0.00442	0.00768	0.01010	0.01274
Kddcup-SVM	0.00112	0.00137	**0.00038**	0.00134	0.00192	0.00250	0.00486	0.00363	0.00154	0.00200
Gisette-LR	0.04637	0.04476	0.03693	0.03683	0.03308	0.03828	0.01420	0.01423	0.02461	0.03560
Gisette-SVM	0.01745	0.03641	0.02911	0.03548	0.04540	0.03603	0.00961	0.00926	0.01188	0.01327

adaptively chosen sequences of queries along with required accuracy bounds. DPella Lobo-Vesga et al. [2020] provides a programming framework for reasoning about privacy, accuracy, and their trade-offs. It uses taint analysis to detect probabilistic independence and derive tighter accuracy bounds using Chernoff bounds. Different from the above studies, our approach focuses on the empirical risk minimization problems and uses empirical error of the learning algorithms to pick the epsilon. Besides, our approach can be directly applied to any objective perturbed learning algorithms.

9 Conclusion

In this paper, we focus on the practical issue of differentially private machine learning: how to assess the impact of different ϵ values on utility of the well-trained model in advance. We formalize this problem as the following question: how ϵ exactly affects the utility. We find that most existing researches of differentially private machine learning refer to this issue. Usually, they define this analysis as "utility analysis" and conclusively give the utility bound based on "utility theorem". However, "utility theorems" often fail to provide meaningful assessments in practice as they tend to be worst-case bounds and would be too conservative as a practical reference.

Different from previous researches, we address this issue through "empirical" theorem. By tracing the change of utility with modification of ϵ in objective perturbation mechanisms, we find that there is an established relationship between ϵ and utility. We reveal this relationship through a formulaic definition and then put forward a practical approach for estimating the utility of the trained model under an arbitrary value of ϵ. Both theoretical analysis and experimental results on real-world datasets demonstrate the good estimation accuracy and broad applicability of our approximation approach. To the best of our knowledge, it is the first time to formalize such a relationship between ϵ and utility under practical cases. The new understanding of utility-privacy trade-off in differential privacy provided by our analysis opens a huge space for exploration in theoretical research and practical application. As providing algorithms with strong utility guarantees that also give privacy when possible becomes more and more accepted, we believe that our approach would have high practical value and can provide meaningful guidance for addressing practical issues.

In our future work, we wish to further improve the generality of our approximation approach, such as extend it to more general perturbation mechanisms (i.e. gradient perturbation) or give further analysis of our approximate approach for general cases (i.e. non-convex, non-convergent and non-differentiable cases).

References

Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Conference on Computer and Communications Security*, pages 1322–1333. ACM, 2015.

Florian Tramèr, Fan Zhang, Ari Juels, Michael K Reiter, and Thomas Ristenpart. Stealing machine learning models via prediction apis. In *USENIX Security Symposium*, pages 601–618. USENIX Association, 2016.

Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *Symposium on Security and Privacy*, pages 3–18. IEEE Computer Society, 2017.

Karan Ganju, Qi Wang, Wei Yang, Carl A Gunter, and Nikita Borisov. Property inference attacks on fully connected neural networks using permutation invariant representations. In *Conference on Computer and Communications Security*, pages 619–633. ACM, 2018.

Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. Privacy risk in machine learning: Analyzing the connection to overfitting. In *Computer Security Foundations Symposium*, pages 268–282. IEEE Computer Society, 2018.

Luca Melis, Congzheng Song, Emiliano De Cristofaro, and Vitaly Shmatikov. Exploiting unintended feature leakage in collaborative learning. In *Symposium on Security and Privacy*, pages 691–706. IEEE, 2019.

L. Sweeney. k-anonymity: a model for protecting privacy. In *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, pages 557–570, 2002.

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

Kun Liu and Evimaria Terzi. Towards identity anonymization on graphs. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, page 93–106. Association for Computing Machinery, 2008.

Ninghui Li, Wahbeh Qardaji, and Dong Su. On sampling, anonymization, and differential privacy or, k-anonymization meets differential privacy. In *Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security*, page 32–33. Association for Computing Machinery, 2012.

Lars Backstrom, Cynthia Dwork, and Jon Kleinberg. Wherefore art thou r3579x? anonymized social networks, hidden patterns, and structural steganography. *Commun. ACM*, 54(12):133–141, 2011.

Arvind Narayanan and Vitaly Shmatikov. De-anonymizing social networks. In *2009 30th IEEE Symposium on Security and Privacy*, pages 173–187, 2009.

Mudhakar Srivatsa and Mike Hicks. Deanonymizing mobility traces: Using social network as a side-channel. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security*, page 628–637. Association for Computing Machinery, 2012.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography*, pages 265–284. Springer Berlin Heidelberg, 2006.

Kamalika Chaudhuri, Claire Monteleoni, and Anand D Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(Mar):1069–1109, 2011.

Daniel Kifer, Adam Smith, and Abhradeep Thakurta. Private convex empirical risk minimization and high-dimensional regression. In *Annual Conference on Learning Theory*, volume 23, pages 25.1–25.40. JMLR.org, 2012.

Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *Annual Symposium on Foundations of Computer Science*, pages 464–473. IEEE, 2014.

Di Wang, Minwei Ye, and Jinhui Xu. Differentially private empirical risk minimization revisited: Faster and more general. In *Advances in Neural Information Processing Systems*, pages 2722–2731, 2017.

Roger Iyengar, Joseph P Near, Dawn Song, Om Thakkar, Abhradeep Thakurta, and Lun Wang. Towards practical differentially private convex optimization. In *Symposium on Security and Privacy*, pages 299–316. IEEE, 2019.

Katrina Ligett, Seth Neel, Aaron Roth, Bo Waggoner, and Steven Z. Wu. Accuracy first: Selecting a differential privacy level for accuracy-constrained erm. In *Advances in Neural Information Processing Systems*, page 2563–2573. Curran Associates Inc., 2017.

Shai Shalev-Shwartz and Nathan Srebro. Svm optimization: Inverse dependence on training set size. In *International Conference on Machine Learning*, page 928–935. ACM, 2008.

Diederik P. Kingma and M. Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2014.

Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, pages 1278–1286. JMLR.org, 2014.

Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *International Conference on Machine Learning*, page 1885–1894. PMLR, 2017.

Jaewoo Lee and Chris Clifton. How much is enough? choosing ϵ for differential privacy. In *International Conference on Information Security*, pages 325–340. Springer, 2011.

Justin Hsu, Marco Gaboardi, Andreas Haeberlen, Sanjeev Khanna, Arjun Narayan, Benjamin C Pierce, and Aaron Roth. Differential privacy: An economic method for choosing epsilon. In *Computer Security Foundations Symposium*, pages 398–410. IEEE Computer Society, 2014.

Maurizio Naldi and Giuseppe Dacquisto. Differential privacy: An estimation theory-based method for choosing epsilon. *arXiv: Cryptography and Security*, 2015.

N. Kohli and P. Laskowski. Epsilon voting: Mechanism design for parameter selection in differential privacy. In *2018 IEEE Symposium on Privacy-Aware Computing (PAC)*, pages 19–30, 2018.

Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.

J. Stamper, A. Niculescu-Mizil, S. Ritter, G.J. Gordon, and K.R. Koedinger. Challenge data set from kdd cup 2010 educational data mining challenge, 2010.

Isabelle Guyon, S R Gunn, Asa Benhur, and Gideon Dror. Result analysis of the nips 2003 feature selection challenge. In *Advances in Neural Information Processing Systems*, pages 545–552, 2004.

Prashanth Mohan, Abhradeep Thakurta, Elaine Shi, Dawn Song, and David Culler. Gupt: Privacy preserving data analysis made easy. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, page 349–360. Association for Computing Machinery, 2012.

Marco Gaboardi, James Honaker, Gary King, Kobbi Nissim, Jonathan R. Ullman, and Salil P. Vadhan. PSI (Ψ): a private data sharing interface. *CoRR*, abs/1609.04340, 2016.

Chang Ge, Xi He, Ihab F. Ilyas, and Ashwin Machanavajjhala. Apex: Accuracy-aware differentially private data exploration. In *Proceedings of the 2019 International Conference on Management of Data*, page 177–194. Association for Computing Machinery, 2019.

Elisabet Lobo-Vesga, Alejandro Russo, and Marco Gaboardi. A programming framework for differential privacy with accuracy concentration bounds. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 411–428, 2020.