Journal of Machine Learning **ISSN**: 2790-2048(e), 2790-203X(p)

Tangent Differential Privacy

Lexing Ying *

Department of Mathematics, Stanford University, Stanford, CA 94305, USA.

Abstract. Differential privacy is a framework for protecting the identity of individual data points in the decision-making process. In this note, we propose a new form of differential privacy, known as tangent differential privacy. Compared to the usual differential privacy, which is defined uniformly across data distributions, tangent differential privacy is tailored to a specific data distribution of interest. It also allows for general distribution distances such as total variation distance and Wasserstein distance. In the context of risk minimization, we demonstrate that entropic regularization ensures tangent differential privacy under relatively general conditions on the risk function.

Keywords:
Differential privacy,
Entropic regularization.

Article Info.:
Volume: 4
Number: 3
Pages: 157 - 165
Date: September/2025
doi.org/10.4208/jml.240928

Article History:
Received: 28/09/2024
Accepted: 12/04/2025

Communicated by:
Song Mei

1 Introduction

Differential privacy is a framework for protecting the identity of individual data points in the machine learning process. The most commonly discussed differential privacy is ϵ-differential privacy. A randomized algorithm is called ϵ-differential private if, for any two input data distributions that differ by one element, the ratio of the probabilities at any outcome is bounded by at most $\exp(\epsilon)$. The definition clearly shows that differential privacy is a uniform concept across all data distributions. In many machine learning applications, one often cares about a specific data distribution and raises privacy concerns about the impact of deleting or adding a single or small number of data points to this specific data distribution.

To address such questions, we propose here tangent differential privacy, a privacy concept tailored to a specific data distribution. When applying the case of risk minimization (such as supervised learning), we demonstrate that entropic regularization guarantees tangent differential privacy under relatively general conditions.

Related work. The concepts of ϵ-differential privacy and (ϵ, δ)-differential privacy were first proposed in [9, 10] and a wonderful resource for this vast literature is [11]. Several efforts have been devoted to relax or reformulate differential privacy, with examples including Renyi differential privacy [17], concentrated differential privacy [6, 12], and Lipschitz privacy [15]. In a broader context, other related forms of privacy concepts have also been developed, such as local differential privacy [8, 13, 14] and the recently proposed metric privacy [4, 5]. The concept of tangent differential privacy proposed here is closely

*Corresponding author. lexing@stanford.edu

related to Lipschitz privacy, though the latter is defined as a uniform concept across all data distributions.

Contents. The rest of the note is organized as follows. Section 2 introduces the concept of tangent differential privacy. Section 3 considers the risk minimization problem and proposes entropic regularization as a solution of tangent differential privacy for both total variation and Wasserstein distances. Section 4 concludes with some discussions.

2 Tangent differential privacy

Let X be the metric space of the data points, and $\mathcal{P}(X)$ be the space of distributions over X. Let W be the metric space of outputs, $\mathcal{P}(W)$ be the space of distributions over W, and $\mathcal{F}(W)$ be the space of bounded functions over W. Here, the output space W can be quite general, such as $\mathbb{R}^n$, the space of regression functions, or the space of neural network weights [1]. To discuss differential privacy, let A be a randomized algorithm that takes $p \in \mathcal{P}(X)$ and produces a randomized output w. Because A is random, we can regard it as a (typically nonlinear) map

$$A : \mathcal{P}(X) \quad \to \quad \mathcal{P}(W),$$

taking $p(x)$ to a distribution $q(w)$. When $q(w)$ has a bounded density, we can also consider its logarithm

$$\log \circ A : \mathcal{P}(X) \quad \to \quad \mathcal{F}(W),$$

taking $p(x)$ to a function $(\log q)(w)$.

Let us denote T_p and T_q as the tangent spaces of signed measures at p and q, respectively. The tangent map of A at p is $DA_p : T_p \to T_q$. Suppose that p is the data distribution of interest. For any $\tilde{p}$ close to p, the linear approximation suggests that

$$A\tilde{p} - Ap \approx DA_p \cdot (\tilde{p} - p). \tag{2.1}$$

In the usual setting, p can be an empirical distribution with N data samples $\{x_i\}$ and $\tilde{p}$ is obtained by removing a distinguished sample x_k

$$p(x) = \frac{1}{N}\sum_{i=1}^{N} \delta_{x_i}(x), \quad \tilde{p}(x) = \frac{1}{N-1}\sum_{i \neq k} \delta_{x_i}(x).$$

This also extends naturally to the situation where $\tilde{p}$ is obtained from p by changing a small number of data points.

Similarly, if $T_{\log q}$ is the tangent space at $\log q$, the tangent map of $\log \circ A$ at p is

$$D(\log \circ A)_p : T_p \quad \to \quad T_{\log q}.$$

For any $\tilde{p}$ close to p, we have

$$\log(A\tilde{p}) - \log(Ap) \approx D(\log \circ A)_p \cdot (\tilde{p} - p). \tag{2.2}$$