# Private Algorithms Can be Always Extended

Christian Borgs* Jennifer Chayes† Adam Smith‡ Ilias Zadik§

October 27, 2018

**Abstract**

We consider the following fundamental question on $\varepsilon$-differential privacy. Consider an arbitrary $\varepsilon$-differentially private algorithm defined on a subset of the input space. Is it possible to extend it to an $\varepsilon'$-differentially private algorithm on the whole input space for some $\varepsilon'$ comparable with $\varepsilon$? In this note we answer affirmatively this question for $\varepsilon' = 2\varepsilon$. Our result applies to every input metric space and space of possible outputs. This result originally appeared in a recent paper by the authors [2]. We present a self-contained version in this note, in the hopes that it will be broadly useful.

## 1 Introduction

We are undoubtedly living in a revolutionary era for data science. The number and magnitude of the available datasets have been growning enormously during the last years. Arguably, though, what makes the data so useful is also what makes them very sensitive. In order to get a theoretical handle on the extent to which an algorithm reveals too much about an individual's data, Dwork et al. [6] introduced *differential privacy*, a mathematical property of statistical algorithms which guarantees privacy of individuals' input data. Roughly speaking, differential privacy requires that a change to one individual's input data not affect the algorithm's output distribution too much.

Differential privacy appears in the literature in various forms. It is typically presented in its "relaxed form" with respect to two parameters $\varepsilon > 0$ and $\delta \geq 0$, [6]. In this note, we focus solely on the special case $\delta = 0$, sometimes called "pure" differential privacy, or just $\varepsilon$-differential privacy. Differential privacy is a property of a randomised algorithm which takes values on some metric space $(\mathcal{M}, d)$ and outputs distributions on some measurable space $(\Omega, \mathcal{F})$:

**Definition 1.1.** *A randomized algorithm $\mathcal{A}$ is $\varepsilon$-differential private if for all subsets $S \in \mathcal{F}$ of the output measurable space $(\Omega, \mathcal{F})$ and data-sets $D_1, D_2$ of the input metric space $(\mathcal{M}, d)$,*

$$\mathbb{P}\left(\mathcal{A}(D_1) \in S\right) \leq \exp\left[\varepsilon d(D_1, D_2)\right] \mathbb{P}\left(\mathcal{A}(D_2) \in S\right). \tag{1.1}$$

The parameter $\varepsilon > 0$ should be treated as a privacy measure; the smaller the $\varepsilon$ the higher the privacy guarantee.

*Microsoft Research New England e-mail: Christian.Borgs@microsoft.com.

†Microsoft Research New England e-mail: jchayes@microsoft.com.

‡Boston University; e-mail: ads22@bu.edu.

§MIT; e-mail: izadik@mit.edu. Research done in part while an intern at Microsoft Research New England.

An intriguing characteristic of analyzing the performance of an $\varepsilon$-differential private algorithms is an inherent trade-off between accuracy and privacy. This is simply the property that increasing the required privacy level of the algorithm (decreasing $\varepsilon > 0$) constrains the accuracy levels of estimation. The natural quantitiative questions rises: how much accuracy is necessarily sacrificed if we entitle our algorithms to a minimum level of privacy? To make the question more precise, suppose that we face an arbitrary statistical estimation question and we restrict ourselves to the use of algorithms of a certain differential privacy level $\varepsilon$. What is the optimal accuracy that can be achieved? How does it compare to the optimal accuracy achieved without any privacy guarantee? Despite the simplicity of this question, in most examples the exact calculation of the underlying rates of estimation remain, to the best of our knowledge, largely open.

A key feature of any estimation question is the generating data distribution. Notably, this distribution does not affect at all the $\varepsilon$-differential privacy constraint, as (1.1) should hold for any pair of input datasets. On the other hand, the assumed generating distribution can massively change the accuracy guarantee. For example, let us consider the following problem, studied in [7, 3, 1]; the analyst receives a sampled graphs $G$ from a distribution over bounded degree undirected graphs on $n$ vertices. The goal is to estimate, in the mean squared error sense, the edge density of the input graph using an $\varepsilon$-differentially private algorithm. As we mentioned above, the algorithm needs to satisfy (1.1) for all pairs of graphs, but it needs to be accurate only on graphs appearing with non-negligible probability as input. In particular, our algorithms suffices to accurately estimate the edge density solely for bounded degree graphs, while for the rest graphs it could potentially be extremely inaccurate. If one knows more about how the graph was generated, then one can further restrict the "interesting" set of inputs—for example, a graph generated from $G(n, p)$ will have cut density approximately $p$ for all cuts, with high probability (for $p$ not too close to zero).

Building on this inherent feature of the problem the following strategy for designing differentially private algorithms has appeared in various forms.

- First design an algorithm $\mathcal{A}_1$ that is both accurate and $\varepsilon$-differentially private on a set $\mathcal{H} \subset \mathcal{M}$ of "typical" elements of the generating data distribution.
- Extend the algorithm $\mathcal{A}_1$ to a private algorithm $\mathcal{A}_2$ on the whole space of inputs $\mathcal{M}$ so that (a) it is $\varepsilon'$-differentially private on the whole space for some $\varepsilon' \geq \varepsilon$ and (b) when the input belongs in $\mathcal{H}$, it outputs the same distribution as the original algorithm $\mathcal{A}_1$.

Various such extension results has appeared in the literature [7, 1, 9, 5, 4] but they are usually tailored for the specific applications each paper discusses. Here, we give a simple proof that such an extension is **always possible** for $\varepsilon' = 2\varepsilon$.

Differential privacy, as defined in Definition 1.1, can be easily shown to be equivalent with the existence of an $\varepsilon$-Lipschitz map from the space of input data $(\mathcal{M}, d)$ to the space of probability measures of some sample space $(\Omega, \mathcal{F})$ endowed with the $\infty$-Renyi divergence metric, $D_\infty(\mu, \nu) = \log \sup_{S \in \mathcal{F}} |\frac{\mu(S)}{\nu(S)}|$. Indeed, the condition (1.1) applied to two datasets $D_1, D_2$ is equivalent with the condition $D_\infty(\mathcal{A}(D_1), \mathcal{A}(D_2)) \leq \varepsilon d(D_1, D_2)$.

Viewed from this perspective, the extension of an $\varepsilon$-differentially private mechanism reduces to a Lipschitz extension question, where a function is assumed to be Lipschitz on a subset of the domain, and the goal is to be extented to a Lipschitz function on the whole domain. The extendability of Lipschitz functions has been thoroughly studied in the field of functional analysis, see for example the Lecture Notes [8]. One standard result in this line of research, states that if the image space is an $\ell_\infty(\Gamma)$ space for some set $\Gamma$, such an extension is always possible with $\varepsilon' = \varepsilon$ (see Theorem

2.1 in [8]). Unfortunately the image of differentiable private algorithms does not have the exact structure of an $\ell_\infty$ space. Despite that, using similar ideas with one of the standard proofs with the $\ell_\infty$ result, we are able to establish the general extension result for $\varepsilon' = 2\varepsilon$.

# 2 The Extension Result

We consider the following statistical model.

### The Model

Let $n \in \mathbb{N}$ and $\varepsilon > 0$. We assume that the analyst's objective is to estimate a certain quantity which belongs in some measurable space $(\Omega, \mathcal{F})$ from input data which takes values in a metric space $(\mathcal{M}, d)$. The analyst is assumed to use for this task a randomized algorithm $\mathcal{A}$ which should be

(1) as highly **accurate** as possible for input data belonging in some *hypothesis set* $\mathcal{H} \subseteq \mathcal{M}$;

(2) $\varepsilon$-**differentially private** for arbitrary pairs of input data-sets from $(\mathcal{M}, d)$.

### Extending Private Algorithms

We now state formally the result described in the note. Consider an arbitrary $\varepsilon$-differentially private algorithm defined on input belonging in some set $\mathcal{H} \subset \mathcal{M}$. We show that it **can be always extended** to a $2\varepsilon$-differentially private algorithm defined for arbitrary input data from $\mathcal{M}$ with the property that if the input data belongs in $\mathcal{H}$, the distribution of output values is the same with the original algorithm. We state formally the result.

**Proposition 2.1** ("Extending Private Algorithms at $\varepsilon$-cost")**.** *Let $\hat{\mathcal{A}}$ be an $\varepsilon$-differentially private algorithm designed for input from $\mathcal{H} \subseteq \mathcal{M}$. Then there exists a randomized algorithm $\mathcal{A}$ defined on the whole input space $\mathcal{M}$ which is $2\varepsilon$-differentially private and satisfies that for every $D \in \mathcal{H}$, $\mathcal{A}(D) \stackrel{d}{=} \hat{\mathcal{A}}(D)$.*

# 3 Applications

In this section, we describe two applications of Theorem 2.1.

### Bounded Degree Graphs

Let $D, n \in \mathbb{N}$ with $D \leq n$. The authors in [7] discuss the following model. Say that an analyst observes a network and wants to build a private algorithm for estimating a graph quantity of the network (e.g. the number of triangles, the degree histogram or the number of edges). Differential privacy here is defined with respect to the node or rewiring distance; the exact definition of this distance is not in the scope of the note but we encourage the interested reader to the discussion in Section 2.1. of [7], or Section 1 in [2]. Now motivated by the apparent sparsity of various real life networks, the authors in [7] face the following question; consider an algorithm which is $\varepsilon$-differentially private over the space of undirected graphs on $n$ vertices with maximum degree at most $D$, which contains all "realistic" graphs for relatively small values of $D$. Can we extend

this algorithm to a differentially private algorithm on the whole space of undirected graphs on $n$ vertices?

Using an efficient approach which is tailored for bounded degree graphs, they establish the existence of an extension algorithm which is $(2\varepsilon, \delta)$-differentially private for some $\delta \geq 0$ (Lemma 6.2 in [7]).

Notice, though, that the question falls exactly into the setting of our note. In particular, for $\mathcal{M}$ the space of undirected graphs on $n$ vertices endowed with the node distance and $\mathcal{H}$ the set of graphs with bounded degree $D$, Theorem 2.1 implies the existence of an $2\varepsilon = (2\varepsilon, 0)$-differentially private extension. Therefore, we conclude the existential aspect of Lemma 6.2. in [7] as a special case of our result, where we improve the $\delta$ to be equal to 0. It is important, though, to notice that unlike the results in [7], our result does not have any efficiency guarantee.

### Estimating Random Graphs

Theorem 2.1 firstly appeared in [2], where the authors discuss the following fundamental and to the best of our knowledge unexplored question: suppose we receive a sample from a $G(n,p)$ Erdos Renyi random graph model, where $n$ is known and $p$ is unknown. How well can we estimate $p$, in the mean squared error sense, using an $\varepsilon$-differentially private algorithm? Note that differential privacy is again understood here with the respect to the node distance. The standard approach for such a question is to add appropriate Laplace noise to the edge density of the observed graph. In [2] it is established that all such estimators can imply at best a rate of the order

$$\frac{1}{n^2} + \frac{1}{n^2\varepsilon^2}.$$

However, it turns out that this rate is suboptimal. The main reason is that for any level of Laplace noise, the estimator can not take into account the "typical" homogeneous structure of an Erdos Renyi graph. In our paper [2] we take advantage of this property and construct a set $\mathcal{H}$ which captures a typical homogeneous structure of an Erdos Renyi graph. We first build an algorithm on graphs belonging in $\mathcal{H}$, apply Theorem 2.1 to extend the algorithm and finally obtain a rate

$$\frac{1}{n^2} + \frac{\log n}{n^3\varepsilon^2}.$$

The authors in [2] provide also a tight lower bound result (Theorem 4.5. in [2]) in the similar case where the graph is sampled from the uniform model $G(n,m)$ and $\varepsilon$ is constant with respect to $n$. The lower bound result proves that the rate obtained using Theorem 2.1 is optimal for the $G(n,m)$ case and suggests the same for $G(n,p)$.

## 4 Proof of Theorem 2.1

We start with a lemma.

**Lemma 4.1.** *Let $\mu$ be a probability measure on $\Omega$ and $\mathcal{A}'$ be a randomized algorithm designed for input from $\mathcal{H}' \subseteq \mathcal{M}$. Suppose that for any $D \in \mathcal{H}'$, $\mathcal{A}'(D)$ is absolutely continuous to $\mu$ and let $f_D$ the Radon-Nikodym derivative $\frac{d\mathcal{A}'(D)}{d\mu}$. Then the following are equivalent*

*(1) $\mathcal{A}'$ is $\varepsilon$-differentially private;*

*(2) For any* $D, D' \in \mathcal{H}$

$$f_{\mathcal{A}'(D)} \leq \exp\left(\varepsilon d(D, D')\right) f_{\mathcal{A}'(D')}, \tag{4.1}$$

*$\mu$-almost surely.*

*Proof.* For the one direction, suppose $\mathcal{A}'$ satisfies (4.1). Then for any set $S \in \mathcal{F}$ we obtain

$$\begin{aligned}\mathbb{P}\left(\mathcal{A}'(D) \in S\right) &= \int_S f_{\mathcal{A}'(D)} d\mu \\ &\leq \exp\left(\varepsilon d(D, D')\right) \int_S f_{\mathcal{A}'(D')} d\mu \\ &= \exp\left(\varepsilon d(D, D')\right) \mathbb{P}\left(\mathcal{A}'(D) \in S\right).\end{aligned}$$

We prove the other direction by contradiction. Consider the set

$$S = \{f_{\mathcal{A}'(D)} > \exp\left(\varepsilon d(D, D')\right) f_{\mathcal{A}'(D')}\} \in \mathcal{F}$$

and assume that $\mu(S) > 0$. By definition on being strictly positive on a set of positive measure

$$\int_S \left[f_{\mathcal{A}'(D)} - \exp\left(\varepsilon d(D, D')\right) f_{\mathcal{A}'(D')}\right] d\mu > 0$$

or equivalently

$$\int_S f_{\mathcal{A}'(D)} d\mu > \exp\left(\varepsilon d(D, D')\right) \int_S f_{\mathcal{A}'(D')} d\mu. \tag{4.2}$$

On the other hand using $\varepsilon$-differential privacy we obtain

$$\begin{aligned}\int_S f_{\mathcal{A}'(D)} d\mu &= \mathbb{P}(\mathcal{A}'(D) \in S) \\ &\leq \exp\left(\varepsilon d(D, D')\right) \mathbb{P}(\mathcal{A}'(D') \in S) \\ &= \exp\left(\varepsilon d(D, D')\right) \int_S f_{\mathcal{A}'(D')} d\mu,\end{aligned}$$

a contradiction with (4.2). This completes the proof of the Lemma. □

Now we establish Theorem 2.1.

*Proof.* Since $\mathcal{H} \neq \varnothing$, let $D_0 \in \mathcal{H}$ and denote by $\mu$ the measure $\hat{\mathcal{A}}(D_0)$. From the definition of differential privacy we know for all $D \in \mathcal{H}$ and $S \in \mathcal{F}$, if $\mathbb{P}\left(\hat{\mathcal{A}}(D_0) \in S\right) = 0$ then $\mathbb{P}\left(\hat{\mathcal{A}}(D) \in S\right) = 0$. In the language of measure theory that means the measure $\hat{\mathcal{A}}(D)$ is absolutely continuous to $\mathcal{A}(D_0)$. By Radon-Nikodym theorem we conclude that there are measurable functions $f_D : \Omega \to [0, +\infty)$ such that for all $S \in \mathcal{F}$,

$$\mathbb{P}\left(\hat{\mathcal{A}}(D) \in S\right) = \int_S f_D d\mu. \tag{4.3}$$

We define now the following randomized algorithm $\mathcal{A}$. For every $D \in \mathcal{M}$, $\mathcal{A}(D)$ samples from $\Omega$ according to the absolutely continuous to $\mu$ distribution with density proportional to

$$\inf_{D' \in \mathcal{H}} \left[\exp\left(\varepsilon d(D, D')\right) f_{\hat{\mathcal{A}}(D')}\right].$$

That is for every $\omega \in \Omega$ its density with respect to $\mu$ is defined as

$$f_{\mathcal{A}(D)}(\omega) = \frac{1}{Z_D} \inf_{D' \in \mathcal{H}} \left[ \exp\left(\varepsilon d(D, D')\right) f_{\hat{\mathcal{A}}(D')}(\omega) \right],$$

where

$$Z_D := \int_\Omega \inf_{D' \in \mathcal{H}} \left[ \left(\varepsilon d(D, D')\right) f_{\hat{\mathcal{A}}(D)'} \right] d\mu.$$

In particular for all $S \in \mathcal{F}$ it holds

$$\mathbb{P}(\mathcal{A}(D) \in S) = \int_S f_{\mathcal{A}(D)} d\mu.$$

We first prove that $\mathcal{A}$ is $2\varepsilon$-differentially private over all pairs of input from $\mathcal{M}$. Using Lemma 4.1 it suffices to prove that for any $D_1, D_2 \in \mathcal{H}$,

$$f_{\mathcal{A}(D_1)} \leq \exp\left(2\varepsilon d(D_1, D_2)\right) f_{\mathcal{A}(D_2)},$$

$\mu$-almost surely. We establish it in particular for every $\omega \in \Omega$. Let $D_1, D_2 \in \mathcal{M}$. Using triangle inequality we obtain for every $\omega \in \Omega$,

$$\begin{aligned}
\inf_{D' \in \mathcal{H}} \left[ \exp\left(\varepsilon d(D_1, D')\right) f_{\hat{\mathcal{A}}(D')}(\omega) \right] &\leq \inf_{D' \in \mathcal{H}} \left[ \exp\left(\varepsilon \left[ d(D_1, D_2) + d(D_2, D') \right]\right) f_{\hat{\mathcal{A}}(D')}(\omega) \right] \\
&= \exp\left(\varepsilon d(D_1, D_2)\right) \inf_{D' \in \mathcal{H}} \left[ \exp\left(\varepsilon d(D, D')\right) f_{\hat{\mathcal{A}}(D')}(\omega) \right],
\end{aligned}$$

which implies that for any $D_1, D_2 \in \mathcal{M}$,

$$\begin{aligned}
Z_{D_1} &= \int_\Omega \inf_{D' \in \mathcal{H}} \left[ \exp\left(\varepsilon d(D_1, D')\right) f_{\hat{\mathcal{A}}(D')} \right] d\mu \\
&\leq \exp\left(\varepsilon d(D_1, D_2)\right) \int_\Omega \inf_{D' \in \mathcal{H}} \left[ \exp\left(\varepsilon d(D_2, D')\right) f_{\hat{\mathcal{A}}(D')}(\omega) \right] d\mu \\
&= \exp\left(\varepsilon d(D_1, D_2)\right) Z_{D_2}.
\end{aligned}$$

Therefore using the above two inequalities we obtain that for any $D_1, D_2 \in \mathcal{H}$ and $\omega \in \Omega$,

$$\begin{aligned}
f_{\mathcal{A}(D_1)}(\omega) &= \frac{1}{Z_{D_1}} \inf_{D' \in \mathcal{H}} \left[ \exp\left(\varepsilon d(D_1, D')\right) f_{\hat{\mathcal{A}}(D')}(\omega) \right] \\
&\leq \frac{1}{\exp\left(-\varepsilon d(D_2, D_1)\right) Z_{D_2}} \exp\left(\varepsilon d(D_1, D_2)\right) \inf_{D' \in \mathcal{H}} \left[ \exp\left(\varepsilon d(D_2, D')\right) f_{\hat{\mathcal{A}}(D')}(\omega) \right] \\
&= \exp\left(2\varepsilon d(D_1, D_2)\right) \frac{1}{Z_{D_2}} \inf_{D' \in \mathcal{H}} \left[ \exp\left(\varepsilon d(D_2, D')\right) f_{\hat{\mathcal{A}}(D')}(\omega) \right] \\
&= \exp\left(2\varepsilon d(D_1, D_2)\right) f_{\mathcal{A}(D_2)}(\omega),
\end{aligned}$$

as we wanted.

Now we prove that for every $D \in \mathcal{H}$, $\mathcal{A}(D) \stackrel{d}{=} \hat{\mathcal{A}}(D)$. Consider an arbitrary $D \in \mathcal{H}$. From Lemma 4.1 we obtain that $\hat{\mathcal{A}}$ is $\varepsilon$-differentially private which implies that for any $D, D' \in \mathcal{H}$

$$f_{\hat{\mathcal{A}}(D)} \leq \exp\left(\varepsilon d(D, D')\right) f_{\hat{\mathcal{A}}(D')}, \tag{4.4}$$

$\mu$-almost surely. Observing that the above inequality holds as $\mu$-almost sure equality if $D' = D$ we obtain that for any $D \in \mathcal{H}$ it holds

$$f_{\hat{\mathcal{A}}(D)}(x) = \inf_{D' \in \mathcal{H}} \left[ \exp\left(\varepsilon d(D, D')\right) f_{\hat{\mathcal{A}}(D')}(x) \right],$$

$\mu$-almost surely. Using that $f_{\hat{\mathcal{A}}(D)}$ is the Radon-Nikodym derivative $\frac{d\hat{\mathcal{A}}(D)}{d\mu}$ we conclude

$$Z_D := \int_\Omega f_{\hat{\mathcal{A}}(D)} d\mu = \mu(\Omega) = 1.$$

Therefore

$$f_{\hat{\mathcal{A}}(D)} = \frac{1}{Z_D} \inf_{D' \in \mathcal{H}} \left[ \exp\left(\varepsilon d(D, D')\right) f_{\hat{\mathcal{A}}(D')} \right],$$

$\mu$-almost surely and hence

$$f_{\hat{\mathcal{A}}(D)} = f_{\mathcal{A}(D)},$$

$\mu$-almost surely. This suffices to conclude that $\hat{\mathcal{A}}(D) \stackrel{d}{=} \mathcal{A}(D)$ as needed.

The proof of Theorem 2.1 is complete. □

## 5 An Open Problem: Efficiency

Theorem 2.1 answers affirmatively the question of extendability of an arbitrary $\varepsilon$-differential private algorithm. An important and interesting open problem for future work is the underlying computational question; under which conditions such an extension can become computationally efficient? This was the focus of much of the existing work [7, 1, 9, 5, 4]; our general result suggests that much greater generality is possible for polynomial-time extensions.

## References


[1] Jeremiah Blocki, Avrim Blum, Anupam Datta, and Or Sheffet. Differentially private data analysis of social networks via restricted sensitivity. In *Innovations in Theoretical Computer Science (ITCS)*, pages 87–96, 2013.

[2] Christian Borgs, Jennifer T. Chayes, Adam D. Smith, and Ilias Zadik. Revealing network structure confidentially: Improved rates for node-private graphon estimation. In *Symposium of the Foundations of Computer Science (FOCS)*, 2018.

[3] Shixi Chen and Shuigeng Zhou. Recursive mechanism: towards node differential privacy and unrestricted joins. In *ACM SIGMOD International Conference on Management of Data*, pages 653–664, 2013.

[4] Rachel Cummings and David Durfee. Individual sensitivity preprocessing for data privacy. *CoRR*, abs/1804.08645, 2018.

[5] Wei-Yen Day, Ninghui Li, and Min Lyu. Publishing graph degree distribution with node differential privacy. In *International Conference on Management of Data SIGMOD*, pages 123–138, 2016.

[6] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference (TCC)*, pages 265–284, 2006.

[7] Shiva Prasad Kasiviswanathan, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Analyzing graphs with node-differential privacy. In *Theory of Cryptography Conference (TCC)*, pages 457–476, 2013.

[8] Assaf Naor. Metric embeddings and lipschitz extensions, lecture notes. 2015.

[9] Sofya Raskhodnikova and Adam D. Smith. Lipschitz extensions for node-private graph statistics and the generalized exponential mechanism. In *Symposium on Foundations of Computer Science (FOCS)*, pages 495–504, 2016.