Benefits and Pitfalls of the Exponential Mechanism with Applications to Hilbert Spaces and Functional PCA

Jordan Awan [1] **Ana Kenney** [1] **Matthew Reimherr** [1] **Aleksandra Slavković** [1]

Abstract

The exponential mechanism is a fundamental tool of Differential Privacy (DP) due to its strong privacy guarantees and flexibility. We study its extension to settings with summaries based on infinite dimensional outputs such as with functional data analysis, shape analysis, and nonparametric statistics. We show that the mechanism must be designed with respect to a specific base measure over the output space, such as a Gaussian process. We provide a positive result that establishes a Central Limit Theorem for the exponential mechanism quite broadly. We also provide a negative result, showing that the magnitude of noise introduced for privacy is asymptotically non-negligible relative to the statistical estimation error. We develop an ϵ-DP mechanism for functional principal component analysis, applicable in separable Hilbert spaces, and demonstrate its performance via simulations and applications to two datasets.

1. Introduction

Data privacy and security have become increasingly critical to society as we continue to collect troves of highly individualized data. In the last decade, we have seen the emergence of new tools and perspectives on data privacy such as *Differential Privacy* (DP), introduced by Dwork et al. (2006), which provides a rigorous and interpretable definition of privacy. Within the DP framework, numerous tools have been developed that achieve DP in a variety of applications and contexts, such as empirical risk minimization (Chaudhuri et al., 2011; Kifer et al., 2012), linear and logistic regression (Chaudhuri & Monteleoni, 2009; Zhang et al., 2012; Yu et al., 2014; Sheffet, 2017; Awan & Slavković, 2018), hypothesis testing (Vu & Slavkovic, 2009; Wang et al., 2015; Gaboardi et al., 2016; Awan & Slavković, 2018; Canonne et al., 2018), network data (Karwa et al., 2016; Karwa & Slavković, 2016), and density estimation (Wasserman & Zhou, 2010), to name a few.

One of the most flexible and convenient DP tools is the *exponential mechanism*, introduced by McSherry & Talwar (2007), which often fits in naturally with estimation techniques from statistics and machine learning. Many estimation procedures can be described as maximizing a particular objective or utility function:

$$\hat{b} = \arg\max_{b \in \mathcal{Y}} \xi(b), \qquad \text{where } \xi : \mathcal{Y} \to \mathbb{R},$$

or, equivalently, minimizing a loss function such as least squares or a negative log-likelihood. The exponential mechanism provides a sanitized version of $\hat{b}$ by using the objective function directly to add noise. The sanitized estimate, $\tilde{b}$, is drawn from a density, $f(b)$, that is proportional to

$$f(b) \propto \exp\left\{\frac{\epsilon}{2\Delta}\xi(b)\right\},$$

where Δ captures the *sensitivity* of the objective function to small perturbations in the data, and ϵ is the desired *privacy budget* (details in Sections 2 and 3). The idea behind this mechanism is to assign higher density values to regions with higher utility. The constant Δ/ϵ adjusts the spread of the density; as the sensitivity increases or as the privacy budget decreases (meaning a decreased disclosure risk), the variability of $\tilde{b}$ increases. A major advantage of such an approach is its use of the objective function from the non-private estimate, $\hat{b}$, which naturally promotes perturbations with higher utility and discourages those with poor utility.

In this paper we study the exponential mechanism, especially as it pertains to functional data analysis (FDA), shape analysis, and nonparametric statistics, where one has a (potentially) infinite dimensional output. Advances in technology and data collection as part of the "big data era" have made such structures more common across a wide variety of fields including economics, genetics, anthropology, and kinesiology, to name a few. For instance, when studying growth trends in children, one can more fully leverage longitudinal information through FDA by treating growth measurements as trajectories or functions rather than using

[1]Department of Statistics, Pennsylvania State University, University Park, Pennsylvania. Correspondence to: Jordan Awan <awan@psu.edu>.

Proceedings of the 36^{th} International Conference on Machine Learning, Long Beach, California, PMLR 97, 2019.

cross-sectional or summary measurements. Such deeply characterized information naturally leads to privacy concerns, though there is currently very little work concerning FDA and statistical data privacy.

We show that the exponential mechanism can be applied in such settings, but requires a specified base measure over the output space $\mathcal{Y}$. We propose using a Gaussian process as the base measure, as these distributions are well studied and easy to implement. We derive a Central Limit Theorem (CLT) for the exponential mechanism quite broadly, meaning we establish asymptotic normality of the mechanism and show it produces $O(1/\sqrt{n})$ noise. However, this result also implies that the magnitude of the noise introduced for privacy is of the same order as the statistical estimation error. In particular, we show that in most natural settings the exponential mechanism does not add an asymptotically negligible noise, even in finite dimensions.

Using our approach, we develop an ϵ-DP mechanism for functional principal component analysis (FPCA), extending the method of Chaudhuri et al. (2013) to separable Hilbert spaces. FPCA is one of the most widely used tools for FDA largely due to the need for dimension reduction when analyzing infinite dimensional data and parameters. Additionally, FPCA characterizes the dominant modes of variation around an overall mean trend function that can be used for exploratory analysis. We show that a Gaussian process base measure enables us to modify the Gibbs sampling procedure of Chaudhuri et al. (2013) to this functional setting. We illustrate the performance of our private FPCA mechansim through simulations, and apply our mechanism to both the Berkeley growth study from the `fda` package (Ramsay et al., 2018) and the Diffusion Tensory Imaging (DTI) dataset from the `refund` package (Goldsmith et al., 2018).

Related Work: This work most directly builds off of Hall et al. (2013) and Mirshani et al. (2017), which develop the first techniques for producing fully functional releases under DP. Another work in this direction is Alda & Rubinstein (2017), in which they use Bernstein polynomial approximations to release functions. Recently, Smith et al. (2018) applied the techniques of Hall et al. (2013) to privatize Gaussian process regression. In their setup, they assume that the predictors are public knowledge, and use this information to carefully tailor the sanitization noise.

There have been a few accuracy bounds regarding exponential mechanism, which can be found in Section 3.4 of Dwork & Roth (2014). However, these results bound the loss in terms of the objective function, rather than in terms of the private release. Wasserman & Zhou (2010) also develop some accuracy bounds for the exponential mechanism, focusing on mean and density estimation. They show that in the mean estimation problem, the exponential mechanism introduces $O(1/\sqrt{n})$ noise. Both Wang et al. (2015) and Foulds et al. (2016) demonstrate the asymptotic normality of the exponential mechanism, when it is of the form of a posterior distribution by using the tools of the Bernstein-von Mises theorem. Our asymptotic analysis of the exponential mechanism agrees in these settings, and extends this result to a large class of objective functions.

Our application to FPCA extends the private PCA method proposed in Chaudhuri et al. (2013). There have been other approaches to private multivariate PCA. Blum et al. (2005) were one of the first to develop a DP procedure for principal components, which is a postprocessing of a noisy covariance matrix. Dwork et al. (2014) follow the same approach and develop bounds for this algorithm; they also develop an online algorithm for private PCA. Jiang et al. (2013) modify this approach by both introducing noise in the covariance matrix as well as to the projection. Imtiaz & Sarwate (2016) also add noise to the covariance matrix, but use a Wishart distribution rather than normal or Laplace noise.

Organization: In Section 2, we review the necessary background of Differential Privacy. In Section 3, we recall the exponential mechanism and give asymptotic results for the performance of the exponential mechanism in both finite and infinite dimensional settings. In Section 4 we show how the exponential mechanism can be applied to produce Functional Principal Components, and in Section 5 we give a Gibbs sampler for this mechanism. In Section 6, we study the performance of the private principal components on both simulated data and on the Berkeley and DTI datasets. Finally, we give our concluding remarks in Section 7. Technical details and proofs are in the Supplementary Material.

2. Differential Privacy

In this section we provide a brief overview of differential privacy (DP). Throughout, we let $\mathcal{X}$ denote an arbitrary set, which represents a particular population, and let $\mathcal{X}^n$ be the n-fold Cartesian product, which represents the collection of all possible samples that could be observed. We begin by defining the *Hamming Distance* between two databases.

Definition 2.1 (Hamming Distance)**.** The bivariate function $\delta : \mathcal{X}^n \times \mathcal{X}^n \to \mathbb{Z}$, which maps $\delta(X, Y) := \#\{i \mid X_i \neq Y_i\}$, is called the *Hamming Distance* on $\mathcal{X}^n$.

It is easy to verify that δ is a metric on $\mathcal{X}^n$. If $\delta(X, Y) = 1$ we call X and Y *adjacent*.

Since we are focused on infinite dimensional objects, we define *Differential Privacy* broadly for any statistical summary. In particular, suppose that $f : \mathcal{X}^n \to \mathcal{Y}$ represents a summary of $\mathcal{X}^n$, and let $\mathscr{F}$ be a σ-algebra of subsets of $\mathcal{Y}$ so that the pair $(\mathcal{Y}, \mathscr{F})$ is a measurable space. From a probabilistic perspective, a *privacy mechanism* is a family of probability measures $\{\mu_X : X \in \mathcal{X}^n\}$ over $\mathcal{Y}$. We can now define what we mean when we say the mechanism sat-

isfies ϵ-DP. While DP was originally introduced in Dwork et al. (2006), Definition 2.2 is similar to the versions given in Wasserman & Zhou (2010) and Kifer & Lin (2010).

Definition 2.2 (Differential Privacy: Dwork et al., 2006)**.** A privacy mechanism $\{\mu_X : X \in \mathcal{X}^n\}$ satisfies ϵ-Differential Privacy (ϵ-DP) if for all $B \in \mathscr{F}$ and adjacent $X, X' \in \mathcal{X}^n$,

$$\mu_X(B) \leq \mu_{X'}(B) \exp(\epsilon).$$

From Definition 2.2, we see that, for an ϵ-DP mechanism, μ_X and $\mu_{X'}$ must be *equivalent measures* (i.e., they agree on sets of measure zero) if $\delta(X, X') = 1$. By transitivity, it follows that μ_X and μ_Y are equivalent measures for any $X, Y \in \mathcal{X}^n$. By the Radon-Nikodym Theorem, we can always therefore interpret DP in terms of densities with respect to a common base measure, ν (if needed, one can always take $\nu = \mu_X$ for an arbitrary $X \in \mathcal{X}^n$).

Proposition 2.3. *Let $\mathscr{M} = \{\mu_X \mid X \in \mathcal{X}^n\}$ be a privacy mechanism over a measurable space $(\mathcal{Y}, \mathscr{F})$. Then $\mathscr{M}$ achieves ϵ-DP if and only if there exists a base measure ν such that $\mu_X \ll \nu$ for all $X \in \mathcal{X}^n$ and the densities $\{f_X : X \in \mathcal{X}^n\}$ (Radon-Nikodym derivatives) of the μ_X (with respect to ν) satisfy*

$$f_X(b) \leq f_{X'}(b) \exp(\epsilon),$$

ν-almost everywhere and for all adjacent $X, X' \in \mathcal{X}^n$.

Proof Sketch. The reverse direction is in Hall et al. (2013). For the other direction, suppose that there exists a set B and adjacent databases X, X' such that $f_X(b) > f_{X'}(b) \exp(\epsilon)$ for all $b \in B$ and that $\nu(B) > 0$. This implies that $\mu_X(B) > \exp(\epsilon)\mu_{X'}(B)$, a contradiction. □

Interpreting DP in terms of densities is common in the DP literature (e.g., Dwork & Roth, 2014; Kifer et al., 2012), however, we could not find a reference for the precise statement and proof, especially for the reverse implication.

3. Exponential Mechanism

One of the earliest mechanisms designed to satisfy ϵ-DP, is the *exponential mechanism*, introduced by McSherry & Talwar (2007). It uses an objective function, which can be the same objective function used for a (non-private) statistical or machine learning analysis, making it especially easy to link DP with existing inferential tools. A simple proof for Proposition 3.1 can be found in McSherry & Talwar (2007).

Proposition 3.1 (Exponential Mechanism: McSherry & Talwar, 2007)**.** *Let $(\mathcal{Y}, \mathscr{F}, \nu)$ be a measure space. Let $\{\xi_X : \mathcal{Y} \to \mathbb{R} \mid X \in \mathcal{X}^n\}$ be a collection of measurable functions. We say that this collection has a finite* sensitivity Δ_ξ*, if*

$$|\xi_X(b) - \xi_{X'}(b)| \leq \Delta_\xi < \infty,$$

for all adjacent X, X' and ν-almost all b. If $\int_{\mathcal{Y}} \exp(\xi_X(b))\, d\nu(b) < \infty$ for all $X \in \mathcal{X}^n$, then the collection of probability measures $\{\mu_X \mid X \in \mathcal{X}^n\}$ with densities f_X (with respect to ν) satisfying

$$f_X(b) \propto \exp\left[\left(\frac{\epsilon}{2\Delta_\xi}\right)\xi_X(b)\right]$$

satisfies ϵ-DP.

We call the set $\{\xi_X \mid X \in \mathcal{X}^n\}$ the *Objective Function*, used in the exponential mechanism. Note that in Proposition 3.1, if ν is a finite measure, $\Delta(\xi) < \infty$, and $\xi_X(b)$ is bounded above for all $X \in \mathcal{X}'$ and ν-almost all b, then one immediately has $\int \exp(\xi_X(b))\, d\nu(b) < \infty$. We will exploit this fact later on as our base measures in infinite dimensions will actually be taken from Gaussian processes, not from any form of Lebesgue measure.

The exponential mechanism offers a general approach to building DP mechanisms, and in fact, a DP mechanism can be expressed as an instantiation of the exponential mechanism, by taking the objective function to be the log-density of the mechanism (McSherry & Talwar, 2007). We remark that the factor of 2 in the exponential mechanism can sometimes be removed (e.g., location families).

Since the solution to many statistical problems can be expressed as the optimizers of some expression, it is natural to set the objective function in the exponential mechanism to this expression. Often such expressions can be expressed as empirical risks, such as the MLE/MAP estimate (Wang et al., 2015), principal component analysis (Chaudhuri et al., 2013), and quantiles of one-dimensional statistics (Smith, 2011). The following result shows that for objective functions of such forms, the noise added by the exponential mechanism is asymptotically normal.

The intuition for the conditions in Theorem 3.2 is to ensure that the objective function has a unique maximizer, can be well approximated by a quadratic form near its maximum, and that the minimizers and objective functions converge to a some well-behaved quantities.

Theorem 3.2 (Utility of Exp Mech)**.** *Assume the observed record, $X_1, \ldots, X_n$, and corresponding sequence of objective functions $\xi_n(b) := \xi_X(b)$, for $b \in \mathbb{R}^p$ satisfy*

1. *$-n^{-1}\xi_n(b)$ are twice differentiable convex functions and there exists a finite $\alpha > 0$ such that the eigenvalues of $-n^{-1}\xi_n(b)''$ are greater than α for all n and $b \in \mathbb{R}^p$;*
2. *the minimizers satisfy $\hat{b} \to b^\star \in \mathbb{R}^p$ and $-n^{-1}\xi_n(\hat{b})'' \to \Sigma^{-1}$ where Σ is a $p \times p$ positive definite matrix;*
3. *ξ_n has finite sensitivity Δ, which is constant in n.*

Assume the base measure has a bounded, differentiable density $g(b)$ which is strictly positive in a neighborhood of $b^\star$. Then the sanitized value $\tilde{b}$ drawn from the exponential mechanism with privacy parameter ϵ is asymptotically normal

$$\sqrt{n}(\tilde{b} - \hat{b}) \xrightarrow{D} N_p\left(0, \left(\frac{2\Delta}{\epsilon}\right)\Sigma\right).$$

Proof Sketch. The proof is based on a second order Taylor expansion of $\xi_X(b)$ about $\hat{b}$. The linear term vanishes, as $\xi'_X(\hat{b}) = 0$, and the error term is $o(1)$. In the limit, we are left with the density of a normal distribution. □

Theorem 3.2 shows that under common conditions (often satisfied by convex empirical risk functions, and log-likelihoods), the noise added by the exponential mechanism is of order $O(1/\sqrt{n})$. We know by the theory of M-estimators and estimating equations (Hardin & Hilbe, 2002) that the non-private solution to the objective functions $\hat{b}$ also converges at rate $O(1/\sqrt{n})$. So, we have that the use of the exponential mechanism in such cases preserves the $1/\sqrt{n}$ convergence rate, but with a sub-optimal asymptotic variance. This means that asymptotically, to achieve the same performance as the non-private estimator, the exponential mechanism requires k times as many samples, where k is some constant larger than 1, depending on ϵ and Δ. However, for many problems, it is possible to construct DP mechanisms that only introduce $O(1/n)$ noise, thus having equivalent asymptotics to the non-private estimator (e.g., Smith, 2011; Awan & Slavković, 2018).

Next, we extend Theorem 3.2 from $\mathbb{R}^p$ to Hilbert spaces. However, we currently only consider base measures that are Gaussian processes.

Theorem 3.3 (Utility of Exp Mech)**.** *Suppose that the observed record, $X_1, \ldots, X_n$, and objective function $\xi_X(b)$, for $b \in \mathcal{H}$ satisfy*

1. *$-n^{-1}\xi_n(b)$ are twice differentiable convex functions and there exists a finite $\alpha > 0$ such that the eigenvalues of $-n^{-1}\xi_n(b)''$ are greater than α for all n and $b \in \mathcal{H}$;*
2. *the minimizers satisfy $\hat{b} \to b^\star \in \mathcal{H}$ and $-n^{-1}\xi_n(\hat{b})''^{-1} \to \Sigma$ where Σ is positive definite nuclear operator (and convergence is wrt this space);*
3. *ξ_n has finite sensitivity Δ, which is constant in n.*

Assume the base measure is taken to be a Gaussian process, $\nu \sim N_{\mathcal{H}}(0, C)$, such that $\Sigma^{-1}C$ is Hilbert-Schmidt, $\Sigma^{-1}C$ is bounded with respect to the Cameron-Martin space (CMS) of C, and $\hat{b}$ lies in the CMS of C for all n. Then the sanitized estimate $\tilde{b}$ is asymptotically normal

$$\sqrt{n}\left(\tilde{b} - \hat{b}\right) \xrightarrow{D} N_{\mathcal{H}}\left(0, \frac{2\Delta}{\epsilon}\Sigma\right).$$

Proof Sketch. The proof strategy is the same as for Theorem 3.3. However, care is taken to ensure that the approximating densities, after the Taylor expansion, are probabilistically equivalent to the base Gaussian process as otherwise the densities are not well defined. Checking limits also becomes more delicate since matrices are replaced with operators. □

Remark 3.4. The requirement that $\Sigma^{-1}C$ is Hilbert-Schmidt can be interpreted as requiring that the base measure be "smoother" than the asymptotic distribution of $\hat{b}$, and ensures the base measure places mass near $\hat{b}$. The second assumption concerning ξ''_X also implies that the sequence of distributions is *tight*. In particular, if one assumed only that Σ was bounded, then the sequence of measures need not be tight and thus one does not get convergence in the "strong topology" in $\mathcal{H}$ (Billingsley, 2013; Chen & White, 1998, Remark 3.3). However, one could still obtain convergence of properly normalized continuous linear functionals.

Example 3.5. Consider $X_1, \ldots, X_n \in \mathcal{H}$ are drawn iid from a Gaussian process with mean μ_X and covariance operator C_X. Consider estimating μ_X using the smooth and convex target function

$$-\xi_X(b) = \sum_{i=1}^{n} \|X_i - b\|^2.$$

Assume that the $\|X_i\| \leq 1$ and thus we need only consider $\|b\| \leq 1$. In that case, the sensitivity is bounded by 4. However, for this target function the exponential mechanism will not be asymptotically Gaussian (in the strong topology). If we consider the second derivative we have $-\xi''_X(b) = 2nI$, and thus $(-\xi''_X(b)/n)^{-1} = (1/2)I$, which is not a nuclear operator in infinite dimensions. However, if instead we consider the penalized version

$$-\xi_X(b) = \sum_{i=1}^{n} \|X_i - b\|^2 + n\lambda\|b\|^2_\Omega,$$

where Ω is a positive-definite nuclear operator and $\|b\|^2_\Omega = \langle b, \Omega^{-1}b\rangle$, then the sensitivity is the same, but the second derivative is now $-\xi''_X(b) = 2nI + 2n\lambda\Omega^{-1}$, whose eigenvalues are bounded from below, as required by Theorem 3.3. In this case,

$$\Sigma = \left(2\lambda\Omega^{-1} + 2I\right)^{-1} = \frac{1}{2}\left(\lambda I + \Omega\right)^{-1}\Omega,$$

which is nuclear as long as Ω is. By choosing a Gaussian process base measure with mean zero, and covariance C such that $\Sigma^{-1}C$ is Hilbert-Schmidt and bounded with respect to the CMS of C, we have by Theorem 3.3 that the noise from the exponential mechanism with privacy parameter ϵ is asymptotically a mean zero Gaussian process with covariance $\frac{8}{\epsilon}\Sigma$.

We stress that, in finite samples, there is no issue regarding privacy even when $\xi_X''(b)^{-1}$ is not nuclear since we are assuming the mechanism is defined using a probability distribution as the base measure. What the previous results and this example illustrate is that there is a price to pay for using such a flexible mechanism. In the "good" case, when the assumptions of Theorem 3.3 are met, one has asymptotically non-negligle noise, but in the "bad" case, the noise can be even larger, since the covariance operator can blow up.

4. DP Functional Principal Components

In this section, we apply the exponential mechanism to the problem of private functional principal component analysis (FPCA). Traditionally, one estimates the principal components, PCs, by first estimating the covariance operator/matrix and then using an eigen decomposition (Kokoszka & Reimherr, 2017). Covariance operators reside in the Hilbert space of Hilbert-Schmidt operators, meaning we can apply Theorem 3.3 to obtain a sanitized covariance estimate and its corresponding utility. Traditional eigenfunction inequalities (Hsing & Eubank, 2015) can then be combined with the post-processing inequality (Dwork & Roth, 2014) to obtain corresponding results for the eigenfunctions. However, we take a different approach here based on Chaudhuri et al. (2013), allowing us to sample the principal component projection more directly. While the privacy guarantees carry over, the utility gains over estimating the covariance directly require a deeper analysis on the manifold of projection operators, which we leave as an open problem.

Let $(\mathcal{H}, \langle\cdot,\cdot\rangle)$ be a Hilbert Space. Let $X \in \mathcal{H}^n$ be such that its components satisfy $\|X_i\| \leq 1$ for all $i = 1, \ldots, n$. Call $\hat{S}(X)$ the k-dimensional subspace of $\mathcal{H}$ given by the span of the first k principal components of X. Let $P_{\hat{S}(X)} : \mathcal{H} \to \mathcal{H}$ denote the projection operator of $\mathcal{H}$ onto $\hat{S}(X)$. We can write $P_{\hat{S}(X)}$ as the solution to the optimization problem

$$P_{\hat{S}(X)} = \arg\min_{P \in \mathscr{P}_k} \sum_{i=1}^{n} \|X_i - PX_i\|^2, \tag{1}$$

where $\mathscr{P}_k$ is the set of projection operators $P : \mathcal{H} \to \mathcal{H}$ of rank k. Equivalently, we can write

$$P_{\hat{S}(X)} = \arg\max_{P \in \mathscr{P}_k} \sum_{i=1}^{n} \|PX_i\|^2.$$

More specifically, we develop a set of probability measures $\mathscr{M}$ on $\mathscr{P}_k$, indexed by $\mathcal{H}^n$, which satisfy ϵ-DP, such that a random element P from $\mu_X \in \mathscr{M}$ is "close" to $P_{\hat{S}(X)}$.

Our approach follows that of Chaudhuri et al. (2013). Our objective function is $\xi : \mathcal{X}^n \times \mathscr{P}_k \to \mathbb{R}$, defined by $\xi_X(P) = \sum_{i=1}^n \|PX_i\|^2$. Note that $\Delta_\xi = 1$, since $\|PX_i\|^2 \leq \|X_i\|^2 \leq 1$ for any $P \in \mathscr{P}_k$ and any $i = 1, \ldots, n$. Since $\sum_{i=1}^n \|PX_i\|^2 \leq n$, for any probability measure ν on $\mathscr{P}_k$, the class of densities on $\mathscr{P}_k$ with respect to ν given by

$$f_X(P) \propto \exp\left(\frac{\epsilon}{2}\sum_{i=1}^{n}\|PX_i\|^2\right), \text{ satisfies } \epsilon\text{-DP.}$$

If $\mathcal{H}$ is finite dimensional, then $\mathscr{P}_k$ is a compact subset of the space of linear operators (e.g. matrices when $\mathcal{H} = \mathbb{R}^p$). In that case, there exists a uniform distribution on $\mathscr{P}_k$. In Chaudhuri et al. (2013), they implement the exponential mechanism as above, with respect to the uniform distribution on $\mathscr{P}_k$.

For arbitrary $\mathcal{H}$, $\mathscr{P}_k$ is not compact, so we must find another base measure on $\mathscr{P}_k$. To understand our proposed construction, we again consider the finite dimensional $\mathcal{H}$. Let $P \sim \text{Unif}(\mathscr{P}_k)$, that is P is drawn from the uniform distribution on $\mathscr{P}_k$. Let $V_1, \ldots, V_k \overset{\text{iid}}{\sim} N(0, I)$, be iid multivariate normal with mean zero and identity covariance matrix. Then $P \overset{d}{=} \text{Projection}(\text{span}(V_1, \ldots, V_k))$ (since V_k is invariant under rotations). From this factorization, a natural extension for arbitrary $\mathcal{H}$ becomes clear. Let $V_1, \ldots, V_k \overset{\text{iid}}{\sim} N_{\mathcal{H}}(0, C)$, be iid Gaussian processes in $\mathcal{H}$ with zero mean and covariance operator C. Note that C must be positive definite and nuclear, which excludes the identity when $\mathcal{H}$ is infinite dimensional (Bogachev, 1998). We can also tailor C to instill certain properties such as smoothness or periodicity. Then set $P = \text{Projection}(\text{span}(V_1, \ldots, V_k))$. This procedure induces in a probability measure on $\mathscr{P}_k$, which we call ν.

Theorem 4.1. *Let $\mathcal{H}$ be a real separable Hilbert Space and $\mathscr{P}_k$ the collection of all k-dimensional projection operators over $\mathcal{H}$. Let ν be the probability measure over $\mathscr{P}_k$ induced by the transformation Projection(span($V_1, \ldots, V_k$)), where $V_i \in \mathcal{H}$ are iid Gaussian process with mean 0 and covariance operator C. Let $X \in \mathcal{H}^n$ be such that its components satisfy $\|X_i\| \leq 1$ for all $i = 1, \ldots, n$. Let $\mathscr{M}$ be the class of probability measures on $\mathscr{P}_k$ with densities*

$$f_X(P) \propto \exp\left(\frac{\epsilon}{2}\sum_{i=1}^{n}\|PX_i\|^2\right)$$

with respect to ν. Then $\mathscr{M}$ satisfies ϵ-DP.

Theorem 4.2. *Let $\mathcal{H}$ be a Hilbert Space, $k < n$ be two positive integers, and $X \in \mathcal{H}^n$ be such that its components satisfy $\|X_i\| \leq 1$ for all $i = 1, \ldots, n$. Define $\mathscr{M}$ as the class of probability measures on $\mathcal{H}^k$ with densities $f_X(V_1, \ldots, V_k)$ proportional to*

$$\exp\left(\frac{\epsilon}{2}\sum_{i=1}^{n}\|\text{Projection}(\text{span}(V_1, \ldots, V_k))X_i\|^2\right)$$

with respect to ν (the measure induced by the Gaussian distribution $N^k(0, C)$) on $\mathcal{H}^k$. Then $\mathscr{M}$ satisfies ϵ-DP.

The proofs of Theorems 4.1 and 4.2 are simple applications of Proposition 3.1, found in the Supplementary Materials.

Theorem 4.2 can be interpreted as outputting an arbitrary basis for a k-dimensional subspace $\widetilde{S}$ of $\mathcal{H}$, which can then be assembled into a projection operator.

Remark 4.3. Often, C can be interpreted as instilling some particular structure on P or the V_i. For example, if $\mathcal{H} = L^2[0,1]$, then C could be defined using the kernel of an RKHS, chosen so that $\widetilde{S}$ have a certain number of derivatives (as many Sobelev spaces are RKHS as well (Berlinet & Thomas-Agnan, 2011)), which is often a natural assumption.

5. PCA continued: Sampling

In the previous section, we developed a set of ϵ-DP probability measures for arbitrary Hilbert Spaces. In this section we specify an efficient method to sample from these distributions. As is common in FDA (Ramsay & Silverman, 2005; Kokoszka & Reimherr, 2017), we use finite dimensional approximations via basis expansions for computation.

Let $b_1, b_2, \ldots$ be an orthonormal basis for $\mathcal{H}$. We will work in the m-dimensional subspace $\mathcal{H}_m = \text{span}(b_1, \ldots, b_m)$. Given our observed values $X_i \in \mathcal{H}$, call $X_{ij} = \langle X_i, b_j \rangle$ for $i = 1, \ldots, n$ and $j = 1, \ldots, m$. Note that this is simply a computational convenience as we have, by assmption, that

$$\|X_i\|^2 = \sum_{i=1}^{\infty} \langle X_i, b_j \rangle^2 = \sum_{i=1}^{m} \langle X_i, b_j \rangle^2 + \sum_{i=m+1}^{\infty} \langle X_i, b_j \rangle^2,$$

meaning that m can be selected such that an overwhelming majority of the variation in the X_i is captured by $b_1, \ldots, b_m$. A subsequent FPCA would then substantially reduce the dimension while controlling information loss. We arrange these real values in an $n \times m$ matrix $\mathbf{X} = (X_{ij})$.

Next, let C be a nuclear covariance operator on $\mathcal{H}$. Write $C_{ij} = \langle b_i, C b_j \rangle$ for $i, j = 1, \ldots, m$. We put these values in an $m \times m$ matrix $\mathbf{C} = (C_{ij})$, which is a positive definite matrix in $\mathbb{R}^{m \times m}$. In this setup, we then draw $(V_1, \ldots, V_k) \in \mathcal{H}_m$. Call $V_{ij} = \langle V_i, b_j \rangle$, and arrange these values into a real-valued matrix $\mathbf{V} = (V_{ij})$ We then draw from the density $f(V)$, with respect to Lebesgue measure on $\mathbb{R}^{k \times m}$, which is proportional to

$$\exp\left(\frac{\epsilon}{2} \text{tr}(X^\top X V (V^\top V)^{-1} V^\top - V^\top C^{-1} V))\right).$$

In fact, we can obtain a more convenient form for sampling. Since we only need the span of V, we can condition on the columns of V being orthonormal. The density $f(V \mid \text{orthonormal})$, with respect to the uniform measure on the set of orthonormal matrices in $\mathbb{R}^{m \times k}$, is proportional to

$$\exp\left(\frac{\epsilon}{2} \text{tr}\left(V^\top \left(X^\top X - C^{-1}\right) V\right)\right),$$

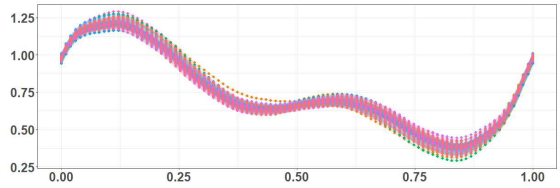

Figure 1: Plot of 100 curves generated for the simulation.

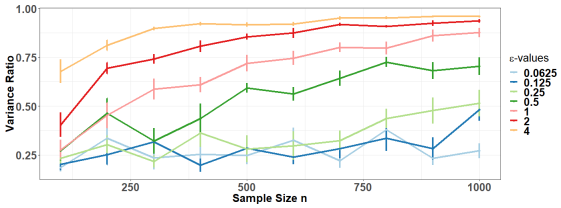

(a) Average ratio of variance explained between the private and non-private principal components.

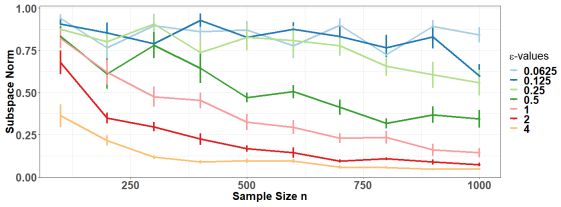

(b) Average subspace norm of private principal components.

Figure 2: Average performance measurements in simulation scenarios over sample sizes ranging from $n = 100$ to 1000. Standard error bars are provided at each point.

which is an instance of the Matrix-Bingham-Von-Mises distribution, for which an efficient Gibbs sampler is known (Hoff, 2009; Hoff & Franks, 2018).

6. Numerical Studies

In this section we assess the numerical performance of the private FPCA method, developed in Sections 4 and 5.

6.1. Simulation Study

For our simulation study, we generated data on a grid of 100 evenly spaced points on $[0,1]$ using the Karhunen-Loeve expansion with Gaussian noise added:

$$X_i(t_{ik}) = \mu(t_{ik}) + \sum_{j=1}^{p} \frac{1}{j^2} U_{ij} u_j(t_{ik}) + \varepsilon_{ik},$$

for $i = 1, \ldots, n$, $k = 1, \ldots, 100$. The $u_j(t)$ are the true functional principal components, ε_{ik} are independent errors sampled from the Gaussian distribution $N(0,1)$, and scores U_{ij} are sampled from $N(0, 0.1)$. Note that for each scenario we re-scale the X_i so that $||X_i||^2 < 1$ for $i = 1, ..., n$.

The $u_j(t)$ are comprised of Fourier basis functions and to

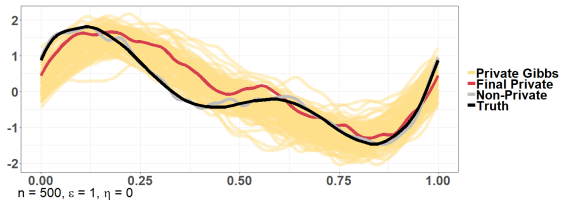

(a) One instance of the simulation where $n = 500$ and $\epsilon = 1$.

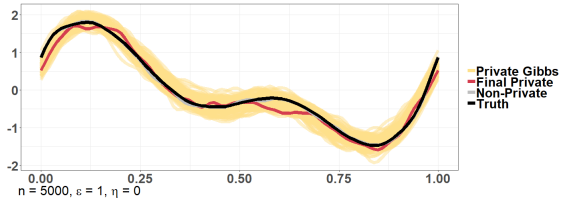

(b) One instance of the simulation where $n = 5000$ and $\epsilon = 1$

Figure 3: Comparisons between the private estimate, non-private estimate, and true first functional principal component. The last 100 Gibbs updates for the private estimate are provided to demonstrate the variability of the mechanism.

fully explore the effectiveness of this approach, we vary the sample size n, privacy budget ϵ, and repeat each scenario 10 times. Data is generated using $p = 21$ true components and additional weights were placed on the fourth term in the Fourier expansion, creating the overall shape shown in Figure 1. We release only $k = 1$ components.

We also specify m, the number of orthonormal basis functions b_i, when restricting the functional observations to a finite dimensional space and C, a nuclear covariance operator on $\mathcal{H}$. It is common to take m to be some sufficiently large value, usually around 40-50, so that results are not sensitive to the truncation. For our simulation scenarios we took $m = 40$ which explained, on average, more than 99% of variation in X. We chose C, to be a diagonal matrix with $C_{ii} = i^{-3}$, which forces the V_i to be continuous. Given that the data is periodic, we use the Fourier basis functions as b_i. Finally, we use the efficient Gibbs sampler of the Matrix-Bingham-Von-Mises distribution (Hoff, 2009), implemented in the `rstiefel` package (Hoff & Franks, 2018) in R. This requires a fixed number of iterations as burn-in prior to starting the procedure. Following the computational experiments in (Chaudhuri et al., 2013), we used 20,000 iterations and had similar convergence results.

We provide two measurements of performance to compare the resulting space of orthogonal projection operators. The first compares the ratio of variability accounted for between the private and non-private estimates of the k functional principal components. More explicitly,

$$0 \leq \frac{||X^T \tilde{P} X||_F^2}{||X^T \hat{P} X||_F^2} \leq 1,$$

where $||\cdot||_F$ is the Frobenius norm, $\tilde{P}$ is the projection onto the span of V drawn from the mechanism in Theorem 4.2, and $\hat{P}$ the the non-private solution to (1).

The second measure gives an indication of how close the range of $\tilde{P}$ is to $\hat{P}$:

$$0 \leq \frac{1}{2}||\tilde{P} - \hat{P}||_F^2 \leq k.$$

If the range of $\tilde{P}$ and $\hat{P}$ agree in h dimensions and are orthogonal in $k - h$ dimensions, then this measure gives the value $k - h$. So this can be interpreted as roughly the number of dimensions that $\tilde{P}$ and $\hat{P}$ disagree.

We summarize the results in Figures 2a and 2b over a range of sample size n and privacy budget ϵ. As expected, larger sample sizes preserve utility (in terms of the two measurements described previously) for stricter privacy requirements. Additionally, we plot a sanitized curve for the first principal component with a sample size of $n = 500$, and $n = 5000$, seen in Figures 3a and 3b. The last 100 Gibbs updates are also shown to demonstrate the variability. Even with a privacy budget of $\epsilon = 1$ and relatively low sample size, the overall shape is captured, but the variance is reduced when $n = 5000$.

6.2. Applications

We applied our private FPCA procedure to two data sets, the Berkeley growth study from the `fda` package (Ramsay et al., 2018), and Diffusion Tensor Imaging (DTI) from the `refund` package (Goldsmith et al., 2018). The Berkeley data has the heights of 93 children at 31 time points with ages between 1-18. DTI gives fractional anisotropy (FA) tract profiles for the corpus callosum (CCA) the right corticospinal tract (RCST) for patients with Multiple Sclerosis and for controls. We study the *cca* data, with 382 patients measured at 93 equally spaced locations of the CCA.

Results are summarized in Tables 1 and 2 when releasing 1-3 principal components across a range of privacy budgets and averaging the performance measurements over 100 repetitions of our procedure. For each data set we selected the Gaussian kernel for C with a smoothness parameter that requires $m = 5$ eigenvalues to explain >99% of variation. Its corresponding eigenfunctions were selected for the orthonormal basis b_i. Our approach is more effective over the DTI data set, which may be due to the true variation explained by the non-private components. For DTI the cumulative variation is .77, .86, and .93 for the top 1, 2, and 3 components respectively, while Berkeley has 0.82, 0.95, and 0.98. When things are too "simple", necessary deviations for privacy show more loss in variation explained compared to the non-private estimates. Overall, this still demonstrates the effectiveness of our procedure under different types of real data with smaller sample sizes.

Table 1: Average performance for private principal components from the Berkeley growth and DTI data sets. Standard errors are provided in parenthesis for reference.

	No. of Components (k)		
	1	2	3
Berkeley	Variance Ratio		
1/8	0.264 (.024)	0.494 (.023)	0.672 (.020)
1/4	0.343 (.024)	0.523 (.023)	0.681 (.020)
1/2	0.408 (.025)	0.523 (.022)	0.729 (.019)
1	0.550 (.025)	0.680 (.018)	0.775 (.015)
2	0.743 (.018)	0.787 (.012)	0.855 (.010)
DTI (*cca*)	Variance Ratio		
1/8	0.372 (.025)	0.569 (.024)	0.727 (.018)
1/4	0.497 (.026)	0.676 (.021)	0.811 (.011)
1/2	0.726 (.020)	0.812 (.014)	0.876 (.009)
1	0.879 (.009)	0.885 (.007)	0.910 (.005)
2	0.933 (.006)	0.928 (.004)	0.939 (.003)

Table 2: Average performance for private principal components from the Berkeley growth and DTI data sets. Standard errors are provided in parenthesis for reference.

	No. of Components (k)		
	1	2	3
Berkeley	Subspace Norm		
1/8	0.776 (.025)	1.115 (.036)	1.100 (.034)
1/4	0.701 (.025)	1.046 (.035)	1.135 (.030)
1/2	0.633 (.027)	1.063 (.033)	1.066 (.030)
1	0.484 (.027)	0.883 (.031)	0.962 (.032)
2	0.275 (.020)	0.770 (.032)	0.938 (.035)
DTI (*cca*)	Subspace Norm		
1/8	0.679 (.026)	1.098 (.035)	1.074 (.030)
1/4	0.544 (.029)	0.976 (.027)	1.079 (.029)
1/2	0.296 (.021)	0.861 (.027)	0.982 (.030)
1	0.131 (.010)	0.770 (.026)	0.940 (.035)
2	0.073 (.006)	0.640 (.030)	0.758 (.035)

7. Discussion

In this paper, we studied the exponential mechanism in the setting of separable Hilbert spaces. We showed that generally when the objective is an empirical risk function, the exponential mechanism has a CLT implying that asymptotically non-negligible noise is introduced. Since the exponential mechanism is popularly used, this result demands the following question: what properties of the objective function guarantee asymptotically negligible noise?

Our asymptotic results extended those in (Wang et al., 2015), which study posterior sampling to achieve DP. In particular, an exponential mechanism can always be viewed as a posterior sample, but often related to a misspecified model. Using this connection, there is a close relationship between Bayesian limit theorems (such as Bernstein-Von Mises) and Theorems 3.2/3.3. However, posterior CLTs often do not hold in arbitrary Hilbert spaces, and the arguments are very delicate (Freedman, 1999). For instance Castillo et al. (2013) avoid densities by working with "nice" projections. Furthermore, posterior CLTs usually require that the likelihood is correct, whereas the exponential mechanism need not correspond to a reasonable model. In contrast, Theorems 3.2/3.3 do not assume any model for the data.

Through our simulations and applications, we found that the choice of C can have a significant impact on the result of the private FPCA analysis. In particular, C can be rescaled by any positive constant, which affects the smoothing but does not change the interpretation in terms of number of derivatives. While our approach requires that C is chosen before seeing the data, it would be preferable to have a method of learning C within the DP procedure. Future researchers should investigate effective methods of tuning parameters under DP.

In the data applications, we found that our DP FPCA approach performs better when there is more variability in the data. This may be because our measures of performance are comparing the DP estimates to the non-private estimates, and the variability hurts both. It would be worth while to investigate this further to better understand how variability in the data affects the performance of DP methods.

As we used a Gibbs sampler to draw approximate samples from our exponential mechanism, it is possible that the Markov chain has not properly mixed, and that the samples are not from the correct distribution. While we know that the Gibbs sampler converges rapidly (Hoff, 2009) and we verified convergence using common heuristics, this is still a potential privacy concern, as the sampling distribution may not satisfy ϵ-DP. This is a problem often encountred when sampling from a non-trivial exponential mechanism. It has been noted that if the sample is drawn from a distribution within a specified total variation of the ϵ-DP distribution, then the sample satisfies (ϵ', δ')-DP for some ϵ' and δ' (see Shen & Yu, 2013, Lemma 5.2 and Wang et al., 2015, Proposition 3). Foulds et al. (2016) provide a different analysis, measuring the privacy cost of Gibbs samplers in particular. Developing rigorous sampling tools for the exponential mechanism is an open problem with on-going research.

Acknowledgments

This research was supported in part by the following grants to Pennsylvania State University: NSF Grant SES-1534433,

NSF Grant DMS-1712826, NIH Grant UL1 TR002014, and NIH Grant 5T32LM012415-03 via the Biomedical Big Data to Knowledge (B2D2K) Predoctoral Training Program. Part of this work was done while the third and fourth authors were visiting the Simons Institute for the Theory of Computing.

References

Alda, F. and Rubinstein, B. I. The bernstein mechanism: Function release under differential privacy. In *AAAI*, pp. 1705–1711, 2017.

Awan, J. and Slavković, A. Differentially private uniformly most powerful tests for binomial data. In *Advances in Neural Information Processing Systems 31*, pp. 4212–4222. Curran Associates, Inc., 2018.

Awan, J. and Slavković, A. Structure and sensitivity in differential privacy: Comparing k-norm mechanisms. *ArXiv e-prints*, January 2018. Submitted.

Berlinet, A. and Thomas-Agnan, C. *Reproducing kernel Hilbert spaces in probability and statistics*. Springer Science & Business Media, 2011.

Billingsley, P. *Convergence of probability measures*. John Wiley & Sons, 2013.

Blum, A., Dwork, C., McSherry, F., and Nissim, K. Practical privacy: the sulq framework. In *Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pp. 128–138. ACM, 2005.

Bogachev, V. I. *Gaussian measures*. Number 62. American Mathematical Soc., 1998.

Canonne, C. L., Kamath, G., McMillan, A., Smith, A. D., and Ullman, J. The structure of optimal private tests for simple hypotheses. *CoRR*, abs/1811.11148, 2018.

Castillo, I., Nickl, R., et al. Nonparametric bernstein–von mises theorems in gaussian white noise. *The Annals of Statistics*, 41(4):1999–2028, 2013.

Chaudhuri, K. and Monteleoni, C. Privacy-preserving logistic regression. In Koller, D., Schuurmans, D., Bengio, Y., and Bottou, L. (eds.), *Advances in Neural Information Processing Systems 21*, pp. 289–296. Curran Associates, Inc., 2009.

Chaudhuri, K., Monteleoni, C., and Sarwate, D. Differentially private empirical risk minimization. In *Journal of Machine Learning Research*, volume 12, pp. 1069–1109, 2011.

Chaudhuri, K., Sarwate, A. D., and Sinha, K. A near-optimal algorithm for differentially-private principal components. *Journal of Machine Learning Research*, 14(1):2905–2943, January 2013. ISSN 1532-4435.

Chen, X. and White, H. Central limit and functional central limit theorems for hilbert-valued dependent heterogeneous arrays with applications. *Econometric Theory*, 14 (2):260–284, 1998.

Dwork, C. and Roth, A. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3–4):211–407, August 2014. ISSN 1551-305X. doi: 10.1561/0400000042.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. *Calibrating Noise to Sensitivity in Private Data Analysis*, pp. 265–284. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006. ISBN 978-3-540-32732-5.

Dwork, C., Talwar, K., Thakurta, A., and Zhang, L. Analyze gauss: Optimal bounds for privacy-preserving principal component analysis. In *Proceedings of the Forty-sixth Annual ACM Symposium on Theory of Computing*, STOC '14, pp. 11–20, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2710-7. doi: 10.1145/2591796.2591883.

Foulds, J., Geumlek, J., Welling, M., and Chaudhuri, K. On the theory and practice of privacy-preserving bayesian data analysis. *arXiv preprint arXiv:1603.07294*, 2016.

Freedman, D. Wald lecture: On the bernstein-von mises theorem with infinite-dimensional parameters. *Ann. Statist.*, 27(4):1119–1141, 08 1999. doi: 10.1214/aos/1017938917.

Gaboardi, M., Lim, H., Rogers, R., and Vadhan, S. Differentially private chi-squared hypothesis testing: Goodness of fit and independence testing. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2111–2120, New York, New York, USA, 20–22 Jun 2016. PMLR.

Goldsmith, J., Scheipl, F., Huang, L., Wrobel, J., Gellar, J., Harezlak, J., McLean, M. W., Swihart, B., Xiao, L., Crainiceanu, C., and Reiss, P. T. *refund: Regression with Functional Data*, 2018. R package version 0.1-17.

Hall, R., Rinaldo, A., and Wasserman, L. Differential privacy for functions and functional data. *Journal of Machine Learning Research*, 14(1):703–727, February 2013. ISSN 1532-4435.

Hardin, J. W. and Hilbe, J. M. *Generalized estimating equations*. Chapman and Hall/CRC, 2002.

Hoff, P. and Franks, A. *rstiefel: Random Orthonormal Matrix Generation and Optimization on the Stiefel Manifold*, 2018. R package version 0.20.

Hoff, P. D. Simulation of the matrix bingham–von mises–fisher distribution, with applications to multivariate and relational data. *Journal of Computational and Graphical Statistics*, 18(2):438–456, 2009.

Hsing, T. and Eubank, R. *Theoretical foundations of functional data analysis, with an introduction to linear operators*. John Wiley & Sons, 2015.

Imtiaz, H. and Sarwate, A. D. Symmetric matrix perturbation for differentially-private principal component analysis. In *Acoustics, Speech and Signal Processing (ICASSP), 2016 IEEE International Conference on*, pp. 2339–2343. IEEE, 2016.

Jiang, X., Ji, Z., Wang, S., Mohammed, N., Cheng, S., and Ohno-Machado, L. Differential-private data publishing through component analysis. *Transactions on data privacy*, 6(1):19, 2013.

Karwa, V. and Slavković, A. Inference using noisy degrees: Differentially private β-model and synthetic graphs. *The Annals of Statistics*, 44(1):87–112, 02 2016. doi: 10.1214/15-AOS1358.

Karwa, V., Krivitsky, P. N., and Slavković, A. B. Sharing social network data: differentially private estimation of exponential family randomgraph models. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 66(3):481–500, 2016. doi: 10.1111/rssc.12185.

Kifer, D. and Lin, B.-R. Towards an axiomatization of statistical privacy and utility. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pp. 147–158. ACM, 2010.

Kifer, D., Smith, A., and Thakurta, A. Private convex empirical risk minimization and high-dimensional regression. *Journal of Machine Learning Research*, 1:1–41, 01 2012.

Kokoszka, P. and Reimherr, M. *Introduction to functional data analysis*. CRC Press, 2017.

McSherry, F. and Talwar, K. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '07, pp. 94–103, Washington, DC, USA, 2007. IEEE Computer Society. ISBN 0-7695-3010-9. doi: 10.1109/FOCS.2007.41.

Mirshani, A., Reimherr, M., and Slavkovic, A. On the Existence of Densities for Functional Data and their Link to Statistical Privacy. *ArXiv e-prints*, November 2017.

Ramsay, J. and Silverman, B. Functional data analysis. Springer, 2005.

Ramsay, J. O., Wickham, H., Graves, S., and Hooker, G. *fda: Functional Data Analysis*, 2018. R package version 2.4.8.

Sheffet, O. Differentially private ordinary least squares. In Precup, D. and Teh, Y. W. (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 3105–3114, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.

Shen, E. and Yu, T. Mining frequent graph patterns with differential privacy. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '13, pp. 545–553, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-2174-7. doi: 10.1145/2487575.2487601.

Smith, A. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC '11, pp. 813–822, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0691-1. doi: 10.1145/1993636.1993743.

Smith, M., lvarez, M., Zwiessele, M., and Lawrence, N. D. Differentially private regression with gaussian processes. In Storkey, A. and Perez-Cruz, F. (eds.), *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pp. 1195–1203, Playa Blanca, Lanzarote, Canary Islands, 09–11 Apr 2018. PMLR.

Vu, D. and Slavkovic, A. Differential privacy for clinical trial data: Preliminary evaluations. In *Proceedings of the 2009 IEEE International Conference on Data Mining Workshops*, ICDMW '09, pp. 138–143, Washington, DC, USA, 2009. IEEE Computer Society. ISBN 978-0-7695-3902-7. doi: 10.1109/ICDMW.2009.52.

Wang, Y.-X., Fienberg, S. E., and Smola, A. J. Privacy for free: Posterior sampling and stochastic gradient monte carlo. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ICML'15, pp. 2493–2502. JMLR.org, 2015.

Wasserman, L. and Zhou, S. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105:489:375–389, 2010.

Yu, F., Rybar, M., Uhler, C., and Fienberg, S. E. Differentially-private logistic regression for detecting multiple-snp association in gwas databases. In *Privacy in Statistical Databases: UNESCO Chair in Data Privacy, International Conference, PSD 2014, Ibiza, Spain,*

September 17-19, 2014. Proceedings, pp. 170–184, Cham, 2014. Springer International Publishing. ISBN 978-3-319-11257-2. doi: 10.1007/978-3-319-11257-2_14.

Zhang, J., Zhang, Z., Xiao, X., Yang, Y., and Winslett, M. Functional mechanism: Regression analysis under differential privacy. *Proc. VLDB Endow.*, 5(11):1364–1375, July 2012. ISSN 2150-8097. doi: 10.14778/2350229.2350253.