
Hermite Polynomial Features for Private Data Generation

Margarita Vinaroz^{*1} Mohammad-Amin Charusaie^{*1} Frederik Harder¹ Kamil Adamczewski¹
Mi Jung Park²

Abstract

Kernel mean embedding is a useful tool to represent and compare probability measures. Despite its usefulness, kernel mean embedding considers infinite-dimensional features, which are challenging to handle in the context of *differentially private* data generation. A recent work (Harder et al., 2021) proposes to approximate the kernel mean embedding of data distribution using *finite-dimensional random features*, which yields analytically tractable sensitivity. However, the number of required random features is excessively high, often ten thousand to a hundred thousand, which worsens the privacy-accuracy trade-off. To improve the trade-off, we propose to replace random features with *Hermite polynomial* features. Unlike the random features, the Hermite polynomial features are *ordered*, where the features at the low orders contain more information on the distribution than those at the high orders. Hence, a relatively low order of Hermite polynomial features can more accurately approximate the mean embedding of the data distribution compared to a significantly higher number of random features. As demonstrated on several tabular and image datasets, Hermite polynomial features seem better suited for private data generation than random Fourier features.

1. Introduction

One of the popular distance metrics for generative modelling is *Maximum Mean Discrepancy* (MMD) (Gretton et al., 2012). MMD computes the average distance between the realizations of two distributions mapped to a reproducing kernel Hilbert space (RKHS). Its popularity is due to several

^{*}Equal contribution ¹Max Planck Institute for Intelligent Systems, Tuebingen, Germany ²University of British Columbia, Vancouver, Canada. CIFAR AI Chair at AMII. Correspondence to: Mi Jung Park <mijungp@cs.ubc.ca>.

facts: (a) MMD can compare two probability measures in terms of all possible moments (i.e., infinite-dimensional features), resulting in no information loss due to a particular selection of moments; and (b) estimating MMD does not require the knowledge of the probability density functions. Rather, MMD estimators are in closed form, which can be computed by pair-wise evaluations of a kernel function using the points drawn from two distributions.

However, using the MMD estimators for training a generator is not well suited when *differential privacy* (DP) of the generated samples is taken into consideration. In fact, the generated points are updated in every training step and the pair-wise evaluations of the kernel function on generated and true data points require accessing data multiple times. One of the key properties of DP is composability that implies each access of data causes privacy loss. Hence, privatizing the MMD estimator in every training step – which is necessary to ensure the resulting generated samples are differentially private – incurs a large privacy loss.

A recent work (Harder et al., 2021), called *DP-MERF*, uses a particular form of MMD via a *random Fourier feature* representation (Rahimi & Recht, 2008) of kernel mean embeddings for DP data generation. Under this representation, one can approximate the MMD in terms of two finite-dimensional mean embeddings (as in eq. 3), where the approximate mean embedding of the true data distribution (data-dependent) is detached from that of the synthetic data distribution (data-independent). Thus, the data-dependent term needs privatization only once and can be re-used repeatedly during training of a generator. However, DP-MERF requires an excessively high number of random features to approximate the mean embedding of data distributions.

We propose to replace¹ the random feature representation of the kernel mean embedding with the *Hermite polynomial* representation. We observe that Hermite polynomial features are ordered where the features at the low orders contain more information on the distribution than those at the high orders. Hence, the required order of Hermite polynomial features is significantly lower than the required number of

¹There are efforts on improving the efficiency of randomized Fourier feature maps, e.g., by using quasi-random points in (Avron et al., 2016).

random features, for the similar quality of the kernel approximation (see Fig. 1). This is useful in reducing the *effective sensitivity* of the data mean embedding. Although the sensitivity is $\frac{1}{m}$ in both cases with the number of data samples m (see Sec. 3), adding noise to a vector of longer length (when using random features) has a worse signal-to-noise ratio, as opposed to adding noise to a vector of shorter length (when using Hermite polynomial features), if we require the norms of these vectors to be the same (for a limited sensitivity). Furthermore, the Hermite polynomial features maintain a better signal-to-noise ratio as it contains more information on the data distribution, even when Hermite polynomial features are the same length as the random Fourier features

To this end, we develop a private data generation paradigm, called *differentially private Hermite polynomials* (DP-HP), which utilizes a novel kernel which we approximate with Hermite polynomial features in the aim of effectively tackling the privacy-accuracy trade-off. In terms of three different metrics we use to quantify the quality of generated samples, our method outperforms the state-of-the-art private data generation methods at the same privacy level. What comes next describes relevant background information before we introduce our method.

2. Background

In the following, we describe the background on kernel mean embeddings and differential privacy.

2.1. Maximum Mean Discrepancy

Given a positive definite kernel $k: \mathcal{X} \times \mathcal{X}$, the MMD between two distributions P, Q is defined as (Gretton et al., 2012): $\text{MMD}^2(P, Q) = \mathbb{E}_{x, x' \sim P} k(x, x') + \mathbb{E}_{y, y' \sim Q} k(y, y') - 2\mathbb{E}_{x \sim P, y \sim Q} k(x, y)$. According to the Moore–Aronszajn theorem (Aronszajn, 1950), there exists a unique reproducing kernel Hilbert space of functions on \mathcal{X} for which k is a reproducing kernel, i.e., $k(x, \cdot) \in \mathcal{H}$ and $f(x) = \langle f, k(x, \cdot) \rangle_{\mathcal{H}}$ for all $x \in \mathcal{X}$ and $f \in \mathcal{H}$, where $\langle \cdot, \cdot \rangle_{\mathcal{H}} = \langle \cdot, \cdot \rangle$ denotes the inner product on \mathcal{H} . Hence, we can find a *feature map*, $\phi: \mathcal{X} \rightarrow \mathcal{H}$ such that $k(x, y) = \langle \phi(x), \phi(y) \rangle_{\mathcal{H}}$, which allows us to rewrite MMD as (Gretton et al., 2012):

$$\text{MMD}^2(P, Q) = \left\| \mathbb{E}_{x \sim P}[\phi(x)] - \mathbb{E}_{y \sim Q}[\phi(y)] \right\|_{\mathcal{H}}^2, \quad (1)$$

where $\mathbb{E}_{x \sim P}[\phi(x)] \in \mathcal{H}$ is known as the (kernel) mean embedding of P , and exists if $\mathbb{E}_{x \sim P} \sqrt{k(x, x)} < \infty$ (Smola et al., 2007). If k is *characteristic* (Sriperumbudur et al., 2011), then $P \mapsto \mathbb{E}_{x \sim P}[\phi(x)]$ is injective, meaning $\text{MMD}(P, Q) = 0$, if and only if $P = Q$. Hence, the MMD associated with a characteristic kernel (e.g., Gaussian kernel) can be interpreted as a distance between the mean embeddings of two distributions.

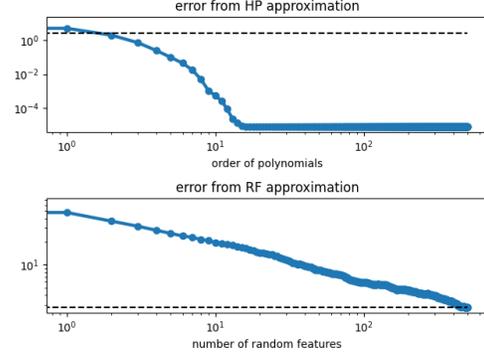


Figure 1. HP VS. RF features. Dataset X contains $N = 100$ samples drawn from $\mathcal{N}(0, 1)$ and X' contains $N = 100$ samples drawn from $\mathcal{N}(1, 1)$. The error is defined by: $\frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N |k(x_i, x'_j) - \hat{\phi}(x_i)^\top \hat{\phi}(x'_j)|$ where $\hat{\phi}$ is either RF or HP features. **Top:** The error decays fast when using HP features (eq. 6). **Bottom:** The plot shows the average error over 100 independent draws of RF features (eq. 4). The error decays slowly when using RF features. The best error (black dotted line) using 500 RF features coincides with the error using HP features with order 2 only.

Given the samples drawn from two distributions: $X_m = \{x_i\}_{i=1}^m \sim P$ and $X'_n = \{x'_j\}_{j=1}^n \sim Q$, we can estimate² the MMD by sample averages (Gretton et al., 2012):

$$\widehat{\text{MMD}}^2(X_m, X'_n) = \frac{1}{m^2} \sum_{i,j=1}^m k(x_i, x_j) + \frac{1}{n^2} \sum_{i,j=1}^n k(x'_i, x'_j) - \frac{2}{mn} \sum_{i=1}^m \sum_{j=1}^n k(x_i, x'_j). \quad (2)$$

However, at $O(mn)$ the computational cost of $\widehat{\text{MMD}}(X_m, X'_n)$ is prohibitive for large-scale datasets.

2.2. Kernel approximation

By approximating the kernel function $k(x, x')$ with an inner product of finite dimensional feature vectors, i.e., $k(x, x') \approx \hat{\phi}(x)^\top \hat{\phi}(x')$ where $\hat{\phi}(x) \in \mathbb{R}^A$ and A is the number of features, the MMD estimator given in eq. 2 can be computed in $O(m + n)$, i.e., linear in the sample size:

$$\widehat{\text{MMD}}^2(P, Q) = \left\| \frac{1}{m} \sum_{i=1}^m \hat{\phi}(x_i) - \frac{1}{n} \sum_{i=1}^n \hat{\phi}(x'_i) \right\|_2^2. \quad (3)$$

This approximation is also beneficial for private data generation: assuming P is a data distribution and Q is a synthetic data distribution, we can summarize the data distribution in terms of its kernel mean embedding (i.e., the first term on the right-hand side of eq. 3), which can be privatized only

²This particular MMD estimator is biased.

once and used repeatedly during training of the generator which produces samples from Q .

2.3. Random Fourier features.

As an example of $\hat{\phi}(\cdot)$, the random Fourier features (Rahimi & Recht, 2008) are derived from the following. Bochner’s theorem (Rudin, 2013) states that for any translation invariant kernel, the kernel can be written as $k(x, x') = \tilde{k}(x - x') = \mathbb{E}_{\omega \sim \Lambda} \cos(\omega^\top (x - x'))$. By drawing random frequencies $\{\omega_i\}_{i=1}^A \sim \Lambda$, where Λ depends on the kernel, (e.g., a Gaussian kernel k corresponds to normal distribution Λ), $\tilde{k}(x - x')$ can be approximated with a Monte Carlo average. The resulting vector of random Fourier features (of length A) is given by

$$\hat{\phi}_{RF, \omega}(x) = (\hat{\phi}_{1, \omega}(x), \dots, \hat{\phi}_{A, \omega}(x))^\top \quad (4)$$

where $\hat{\phi}_{j, \omega}(x) = \sqrt{2/A} \cos(\omega_j^\top x)$, $\hat{\phi}_{j+A/2, \omega}(x) = \sqrt{2/A} \sin(\omega_j^\top x)$, for $j = 1, \dots, A/2$.

DP-MERF (Harder et al., 2021) uses this very representation of the feature map given in eq. 4, and minimizes eq. 3 with a privatized data mean embedding to train a generator.

2.4. Hermite polynomial features.

For another example of $\hat{\phi}(\cdot)$, one could also start with the Mercer’s theorem (See Appendix Sec. C), which allows us to express a positive definite kernel k in terms of the eigen-values λ_i and eigen-functions f_i : $k(x, x') = \sum_{i=1}^{\infty} \lambda_i f_i(x) f_i^*(x')$, where $\lambda_i > 0$ and complex conjugate is denoted by $*$. The resulting *finite-dimensional* feature vector is simply $\hat{\phi}(x) = \hat{\phi}_{HP}(x) = [\sqrt{\lambda_0} f_0(x), \sqrt{\lambda_1} f_1(x), \dots, \sqrt{\lambda_C} f_C(x)]$, where the cut-off is made at the C -th eigen-value and eigen-function. For the commonly-used Gaussian kernel, $k(x, x') = \exp(-\frac{1}{2l^2}(x - x')^2)$, where l is the length scale parameter, an analytic form of eigen-values and eigen-functions are available, where the eigen-functions are represented with Hermite polynomials (See Sec. 3 for definition). This is the approximation we will use in our method.

2.5. Differential privacy

Given privacy parameters $\epsilon \geq 0$ and $\delta \geq 0$, a mechanism \mathcal{M} is (ϵ, δ) -DP if the following equation holds: $\Pr[\mathcal{M}(\mathcal{D}) \in S] \leq e^\epsilon \cdot \Pr[\mathcal{M}(\mathcal{D}') \in S] + \delta$, for all possible sets of the mechanism’s outputs S and all neighbouring datasets \mathcal{D} , \mathcal{D}' differing by a single entry. In this paper, we use the *Gaussian mechanism* to ensure the output of our algorithm is DP. Consider a function $h : \mathcal{D} \mapsto \mathbb{R}^p$, where we add noise for privacy and the level of noise is calibrated to the *global sensitivity* (Dwork et al., 2006), Δ_h , defined by the maximum difference in terms of L_2 -norm $\|h(\mathcal{D}) - h(\mathcal{D}')\|_2$, for neighbouring \mathcal{D} and \mathcal{D}' (i.e. \mathcal{D} and \mathcal{D}' have one sample

difference by replacement). where the output is denoted by $\tilde{h}(\mathcal{D}) = h(\mathcal{D}) + n$, where $n \sim \mathcal{N}(0, \sigma^2 \Delta_h^2 \mathbf{I}_p)$. The perturbed function $\tilde{h}(\mathcal{D})$ is (ϵ, δ) -DP, where σ is a function of ϵ and δ and can be numerically computed using, e.g., the auto-dp package by (Wang et al., 2019).

3. Our method: DP-HP

3.1. Approximating the Gaussian kernel using Hermite polynomials (HP)

Using the *Mehler formula*³ (Mehler, 1866), for $|\rho| < 1$, we can write down the Gaussian kernel⁴ as a weighted sum of Hermite polynomials

$$\exp\left(-\frac{\rho}{1-\rho^2}(x-y)^2\right) = \sum_{c=0}^{\infty} \lambda_c f_c(x) f_c(y) \quad (5)$$

where the c -th eigen-value is $\lambda_c = (1-\rho)\rho^c$ and the c -th eigen-function is defined by f_c , where $f_c(x) = \frac{1}{\sqrt{N_c}} H_c(x) \exp\left(-\frac{\rho}{1+\rho} x^2\right)$, and $N_c = 2^c c! \sqrt{\frac{1-\rho}{1+\rho}}$. Here, $H_c(x) = (-1)^c \exp(x^2) \frac{d^c}{dx^c} \exp(-x^2)$ is the c -th order physicist’s Hermite polynomial.

As a result of the Mehler formula, we can define a C -th order Hermite polynomial features as a feature map (a vector of length $C+1$):

$$\hat{\phi}_{HP}^{(C)}(x) = [\sqrt{\lambda_0} f_0(x), \dots, \sqrt{\lambda_C} f_C(x)], \quad (6)$$

and approximate the Gaussian kernel via $\exp\left(-\frac{\rho}{1-\rho^2}(x-y)^2\right) \approx \hat{\phi}_{HP}^{(C)}(x)^\top \hat{\phi}_{HP}^{(C)}(y)$.

This feature map provides us with a uniform approximation to the MMD in eq. 1, for every pair of distributions P and Q (see Theorem C.1 and Lemma C.1 in Appendix Sec. C).

We compare the accuracy of this approximation with random features in Fig. 1, where we fix the length scale to the median heuristic value⁵ in both cases. Note that the bottom plot shows the average error across 100 independent draws of random Fourier features. We observe that the error decay is significantly faster when using HPs than using RFs. For completeness, we derive the kernel approximation error under HP features and random features for 1-dimensional data in Appendix Sec. B. Additionally, we visualize the effect of length scale on the error further in Appendix Sec. A.

Computing the Hermite polynomial features. Hermite polynomials follow the recursive definition: $H_{c+1}(x) =$

³This formula can be also derived from the Mercer’s theorem as shown in (Zhu et al., 1997; Rasmussen & Williams, 2005).

⁴The length scale l in terms of ρ is $\frac{1}{2l^2} = \frac{\rho}{1-\rho^2}$.

⁵Median heuristic is a commonly-used heuristic to choose a length scale, which picks a value in the middle range (i.e., median) of $\|x_i - x_j\|$ for $1 \leq i, j \leq n$ for the dataset of n samples.

$2xH_c(x) - 2cH_{c-1}(x)$. At high orders, the polynomials take on large values, leading to numerical instability. So we compute the re-scaled term $\phi_c = \sqrt{\lambda_c} f_c$ iteratively using a similar recursive expression given in Appendix Sec. E.

3.2. Handling multi-dimensional inputs

3.2.1. TENSOR (OR OUTER) PRODUCT KERNEL

The Mehler formula holds for 1-dimensional input space. For D -dimensional inputs $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^D$, where $\mathbf{x} = [x_1, \dots, x_D]$ and $\mathbf{x}' = [x'_1, \dots, x'_D]$, the *generalized Hermite Polynomials* (Proposition C.3 and Remark 1 in Appendix Sec. C) allows us to represent the multivariate Gaussian kernel $k(\mathbf{x}, \mathbf{x}')$ by a tensor (or outer) products of the Gaussian kernel defined on each input dimension, where the coordinate-wise Gaussian kernel is approximated with Hermite polynomials:

$$\begin{aligned} k(\mathbf{x}, \mathbf{x}') &= k_{X_1} \otimes k_{X_2} \cdots \otimes k_{X_D} = \prod_{d=1}^D k_{X_d}(x_d, x'_d), \\ &\approx \prod_{d=1}^D \hat{\phi}_{HP,d}^{(C)}(x_d)^\top \hat{\phi}_{HP,d}^{(C)}(x'_d), \end{aligned} \quad (7)$$

where $\hat{\phi}_{HP,d}^{(C)}(\cdot)$ ⁶ is defined in eq. 6. The corresponding feature map, from $k(\mathbf{x}, \mathbf{x}') \approx \mathbf{h}_p(\mathbf{x})^\top \mathbf{h}_p(\mathbf{x}')$, is written as

$$\begin{aligned} \mathbf{h}_p(\mathbf{x}) &= \text{vec} \left[\hat{\phi}_{HP,1}^{(C)}(x_1) \otimes \hat{\phi}_{HP,2}^{(C)}(x_2) \otimes \cdots \otimes \hat{\phi}_{HP,D}^{(C)}(x_D) \right] \end{aligned} \quad (8)$$

where \otimes denotes the tensor (outer) product and vec is an operation that vectorizes a tensor. The size of the feature map is $(C+1)^D$, where D is the input dimension of the data and C is the chosen order of the Hermite polynomials. This is prohibitive for the datasets we often deal with, e.g., for MNIST ($D = 784$) with a relatively small order (say $C = 10$), the size of feature map is 11^{784} , impossible to fit in a typical size of memory.

In order to handle high-dimensional data in a computationally feasible manner, we propose the following approximation. First we subsample input dimensions where the size of the selected input dimensions is denoted by D_{prod} . We then compute the feature map only on those selected input dimensions denoted by $\mathbf{x}^{D_{prod}}$. We repeat these two steps during training. The size of the feature map becomes $(C+1)^{D_{prod}}$, significantly lower than $(C+1)^D$ if $D_{prod} \ll D$. What we lose in return is the injectivity of the Gaussian kernel on the full input distribution, as we compare two distributions

⁶One can let each coordinate's Hermite Polynomials $\hat{\phi}_{HP,d}^{(C)}(x_d)$ take different values of ρ , which determine a different level of fall-offs of the eigen-values and a different range of values of the eigen-functions. Imposing a different cut-off C for each coordinate is also possible.

in terms of selected input dimensions. We need a quantity that is more computationally tractable and also helps distinguishing two distributions, which we describe next.

3.2.2. SUM KERNEL

Here, we define another kernel on the joint distribution over (x_1, \dots, x_D) . The following kernel is formed by defining a 1-dimensional Gaussian kernel on each of the input dimensions:

$$\begin{aligned} \tilde{k}(\mathbf{x}, \mathbf{x}') &= \frac{1}{D} [k_{X_1}(x_1, x'_1) + \cdots + k_{X_D}(x_D, x'_D)], \\ &= \frac{1}{D} \sum_{d=1}^D k_{X_d}(x_d, x'_d), \\ &\approx \frac{1}{D} \sum_{d=1}^D \hat{\phi}_{HP,d}^{(C)}(x_d)^\top \hat{\phi}_{HP,d}^{(C)}(x'_d), \end{aligned} \quad (9)$$

where $\hat{\phi}_{HP,d}^{(C)}(\cdot)$ is given in eq. 6. The corresponding feature map, from $\tilde{k}(\mathbf{x}, \mathbf{x}') \approx \mathbf{h}_s(\mathbf{x})^\top \mathbf{h}_s(\mathbf{x}')$, is represented by

$$\mathbf{h}_s(\mathbf{x}) = \begin{bmatrix} \hat{\phi}_{HP,1}^{(C)}(x_1)/\sqrt{D} \\ \hat{\phi}_{HP,2}^{(C)}(x_2)/\sqrt{D} \\ \vdots \\ \hat{\phi}_{HP,D}^{(C)}(x_D)/\sqrt{D} \end{bmatrix} \in \mathbb{R}^{((C+1) \cdot D) \times 1}, \quad (10)$$

where the features map is the size of $(C+1)D$. For the MNIST digit data ($D = 784$), with a relatively small order, say $C = 10$, the size of the feature map is $11 \times 784 = 8624$ dimensional, which is manageable compared to the size (11^{784}) of the feature map under the generalized Hermite polynomials.

Note that the sum kernel does not approximate the Gaussian kernel defined on the joint distribution over all the input dimensions. Rather, the assigned Gaussian kernel *on each dimension is characteristic*. The Lemma D.1 in Appendix Sec. D shows that by minimizing the approximate MMD between the real and synthetic data distributions based on feature maps given in eq. 10, we assure that the marginal probability distributions of the synthetic data converges to those of the real data.

3.2.3. COMBINED KERNEL

Finally we arrive at a new kernel, which comes from a sum of the two fore-mentioned kernels:

$$k_c(\mathbf{x}, \mathbf{x}') = k(\mathbf{x}, \mathbf{x}') + \tilde{k}(\mathbf{x}, \mathbf{x}'), \quad (11)$$

where $k(\mathbf{x}, \mathbf{x}') \approx \mathbf{h}_p(\mathbf{x}^{D_{prod}})^\top \mathbf{h}_p(\mathbf{x}'^{D_{prod}})$ and $\tilde{k}(\mathbf{x}, \mathbf{x}') \approx \mathbf{h}_s(\mathbf{x})^\top \mathbf{h}_s(\mathbf{x}')$, and consequently the corresponding feature map is given by

$$\mathbf{h}_c(\mathbf{x}) = \begin{bmatrix} \mathbf{h}_p(\mathbf{x}^{D_{prod}}) \\ \mathbf{h}_s(\mathbf{x}) \end{bmatrix} \quad (12)$$

where the size of the feature map is $\mathbb{R}^{((C+1)^{D_{prod}}+(C+1)\cdot D)\times 1}$.

Why this kernel? When D_{prod} goes to D , the product kernel itself in eq. 11 becomes characteristic, which allows us to reliably compare two distributions. However, for computational tractability, we are restricted to choose a relatively small D_{prod} to subsample the input dimensions, which forces us to lose information on the distribution over the un-selected input dimensions. The use of sum kernel is to provide extra information on the un-selected input dimensions at a particular training step. Under our kernel in eq. 11, every input dimension's marginal distributions are compared between two distributions in all the training steps due to the sum kernel, while some of the input dimensions are chosen to be considered for more detailed comparison (e.g., high-order correlations between selected input dimensions) due to the outer product kernel.

3.3. Approximate MMD for classification

For classification tasks, we define a mean embedding for the joint distribution over the input and output pairs (\mathbf{x}, \mathbf{y}) , with the particular feature map given by \mathbf{g}

$$\hat{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}}(\mathcal{D}) = \frac{1}{m} \sum_{i=1}^m \mathbf{g}(\mathbf{x}_i, \mathbf{y}_i). \quad (13)$$

Here, we define the feature map as an outer product between the input features represented by eq. 12 and the output labels represented by one-hot-encoding $\mathbf{f}(\mathbf{y}_i)$:

$$\mathbf{g}(\mathbf{x}_i, \mathbf{y}_i) = \mathbf{h}_c(\mathbf{x}_i) \mathbf{f}(\mathbf{y}_i)^T. \quad (14)$$

Given eq. 14, we further decompose eq. 13 into two, where the first term corresponds to the outer product kernel denoted by $\hat{\boldsymbol{\mu}}_P^p$ and the second term corresponds to the sum kernel denoted by $\hat{\boldsymbol{\mu}}_P^s$:

$$\hat{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}} = \begin{bmatrix} \hat{\boldsymbol{\mu}}_P^p \\ \hat{\boldsymbol{\mu}}_P^s \end{bmatrix} = \begin{bmatrix} \frac{1}{m} \sum_{i=1}^m \mathbf{h}_p(\mathbf{x}_i^{D_{prod}}) \mathbf{f}(\mathbf{y}_i)^T \\ \frac{1}{m} \sum_{i=1}^m \mathbf{h}_s(\mathbf{x}_i) \mathbf{f}(\mathbf{y}_i)^T \end{bmatrix}. \quad (15)$$

Similarly, we define an approximate mean embedding of the synthetic data distribution by $\hat{\boldsymbol{\mu}}_{Q_{\mathbf{x}',\mathbf{y}'}}(\mathcal{D}'_{\boldsymbol{\theta}}) = \frac{1}{n} \sum_{i=1}^n \mathbf{g}(\mathbf{x}'_i(\boldsymbol{\theta}), \mathbf{y}'_i(\boldsymbol{\theta}))$, where $\boldsymbol{\theta}$ denotes the parameters of a synthetic data generator. Then, the approximate MMD is given by: $\widehat{\text{MMD}}_{HP}^2(P, Q) = \|\hat{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}}(\mathcal{D}) - \hat{\boldsymbol{\mu}}_{Q_{\mathbf{x}',\mathbf{y}'}}(\mathcal{D}'_{\boldsymbol{\theta}})\|_2^2 = \|\hat{\boldsymbol{\mu}}_P^p - \hat{\boldsymbol{\mu}}_{Q_{\boldsymbol{\theta}}}^p\|_2^2 + \|\hat{\boldsymbol{\mu}}_P^s - \hat{\boldsymbol{\mu}}_{Q_{\boldsymbol{\theta}}}^s\|_2^2$. In practice, we minimize the augmented approximate MMD:

$$\min_{\boldsymbol{\theta}} \gamma \|\hat{\boldsymbol{\mu}}_P^p - \hat{\boldsymbol{\mu}}_{Q_{\boldsymbol{\theta}}}^p\|_2^2 + \|\hat{\boldsymbol{\mu}}_P^s - \hat{\boldsymbol{\mu}}_{Q_{\boldsymbol{\theta}}}^s\|_2^2. \quad (16)$$

where γ is a positive constant (a hyperparameter) that helps us to deal with the scale difference in the two terms (depending on the selected HP orders and subsampled input

dimensions) and also allows us to give a different importance on one of the two terms. We provide the details on how γ plays a role and whether the algorithm is sensitive to γ in Sec. 5. Minimizing eq. 16 yields a synthetic data distribution over the input and output, which minimizes the discrepancy in terms of the particular feature map eq. 15 between synthetic and real data distributions.

3.4. Differentially private data samples

For obtaining privacy-preserving synthetic data, all we need to do is privatizing $\hat{\boldsymbol{\mu}}_P^p$ and $\hat{\boldsymbol{\mu}}_P^s$ given in eq. 15, then training a generator. We use the Gaussian mechanism to privatize both terms. See Appendix Sec. F for sensitivity analysis. Unlike $\hat{\boldsymbol{\mu}}_P^s$ that can be privatized only and for all, we need to privatize $\hat{\boldsymbol{\mu}}_P^p$ every time we redraw the subsampled input dimensions. We split a target ϵ into two such that $\epsilon = \epsilon_1 + \epsilon_2$ (also the same for δ), where ϵ_1 is used for privatizing $\hat{\boldsymbol{\mu}}_P^p$ and ϵ_2 is used for privatizing $\hat{\boldsymbol{\mu}}_P^s$. We further compose the privacy loss incurred in privatizing $\hat{\boldsymbol{\mu}}_P^p$ during training by the analytic moments accountant (Wang et al., 2019), which returns the privacy parameter σ as a function of (ϵ_2, δ_2) . In the experiments, we subsample the input dimensions for the outer product kernel in every epoch as opposed to in every training step for an economical use of ϵ_2 .

4. Related Work

Approaches to differentially private data release can be broadly sorted into three categories. One line of prior work with background in learning theory aims to provide theoretical guarantees on the utility of released data (Snoke & Slavković, 2018; Mohammed et al., 2011; Xiao et al., 2010; Hardt et al., 2012; Zhu et al., 2017). This usually requires strong constraints on the type of data and the intended use of the released data.

A second line of work focuses on the sub-problem of discrete data with limited domain size, which is relevant to tabular datasets (Zhang et al., 2017; Qardaji et al., 2014; Chen et al., 2015; Zhang et al., 2021). Such approaches typically approximate the structure in the data by identifying small sub-sets of features with high correlation and releasing these lower order marginals in a private way. Some of these methods have also been successful in the recent NIST 2018 Differential Privacy Synthetic Data Challenge (nis), while these methods often require discretization of the data and do not scale to higher dimensionality in arbitrary domains.

The third line of work aims for broad applicability without constraints on the type of data or the kind of downstream tasks to be used. Recent approaches attempt to leverage the modeling power of deep generative models in the private setting. While work on VAEs exists (Acs et al., 2018), GANs are the most popular model (Xie et al., 2018; Torkzadehma-

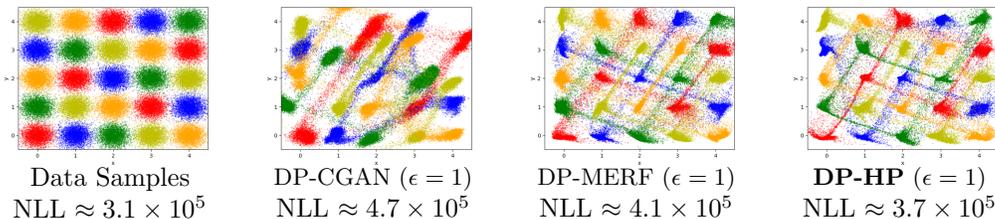


Figure 2. Simulated example from a Gaussian mixture. **Left:** Data samples drawn from a Gaussian Mixture distribution with 5 classes (each color represents a class). NLL denotes the negative log likelihood of the samples given the true data distribution. **Middle-Left:** Synthetic data generated by DP-CGANs at $\epsilon = 1$, where some modes are dropped, which is reflected in poor NLL. **Middle-Right:** Synthetic data samples generated by DP-MERF at $\epsilon = 1$. **Right:** Synthetic data samples generated by DP-HP at $\epsilon = 1$. Our method captures all modes accurately at $\epsilon = 1$, and achieves better NLL thanks to a smaller size of feature map than that of DP-MERF (see text).

hani et al., 2019; Frigerio et al., 2019; Yoon et al., 2019; Chen et al., 2020), where most of these utilize a version of DP-SGD (Abadi et al., 2016) to accomplish this training, while PATE-GAN is based on the private aggregation of teacher ensembles (PATE) (Papernot et al., 2017).

The closest prior work to the proposed method is DP-MERF (Harder et al., 2021), where kernel mean embeddings are approximated with random Fourier features (Rahimi & Recht, 2008) instead of Hermite polynomials. Random feature approximations of MMD have also been used with DP (Balog et al., 2018; Sarpatwar et al., 2019). A recent work utilizes the Sinkhorn divergence for private data generation (Cao et al., 2021), which more or less matches the results of DP-MERF when the regularizer is large and the cost function is the L2 distance. To our knowledge, ours is the first work using Hermite polynomials to approximate MMD in the context of differentially private data generation.

5. Experiments

Here, we show the performance of our method tested on several real world datasets. Evaluating the quality of generated data itself is challenging. Popular metrics such as inception score and Fréchet inception distance are appropriate to use for evaluating color images. For the generated samples for tabular data and black and white images, we use the following three metrics: (a) Negative log-likelihood of generated samples given a ground truth model in Sec. 5.1; (b) α -way marginals of generated samples in Sec. 5.2 to judge whether the generated samples contain a similar correlation structure to the real data; (c) Test accuracy on the real data given classifiers trained with generated samples in Sec. 5.3 to judge the generalization performance from synthetic to real data.

As comparison methods, we tested PrivBayes (Zhang et al., 2017), DP-CGAN (Torkzadehmahani et al., 2019), DP-GAN (Xie et al., 2018) and DP-MERF (Harder et al., 2021). For image datasets we also trained GS-WGAN (Chen et al., 2020). Our experiments were implemented

in PyTorch (Paszke et al., 2019) and run using Nvidia Kepler20 and Kepler80 GPUs. Our code is available at <https://github.com/ParkLabML/DP-HP>.

5.1. 2D Gaussian mixtures

We begin our experiments on Gaussian mixtures, as shown in Fig. 2 (left). We generate 4000 samples from each Gaussian, reserving 10% for the test set, which yields 90000 training samples from the following distribution: $p(\mathbf{x}, \mathbf{y}) = \prod_i^N \sum_{j \in C_{\mathbf{y}_i}} \frac{1}{C} \mathcal{N}(\mathbf{x}_i | \boldsymbol{\mu}_j, \sigma \mathbf{I}_2)$ where $N = 90000$, and $\sigma = 0.2$. $C = 25$ is the number of clusters and C_y denotes the set of indices for means $\boldsymbol{\mu}$ assigned to class y . Five Gaussians are assigned to each class, which leads to a uniform distribution over \mathbf{y} and 18000 samples per class. We use the negative log likelihood (NLL) of the samples under the true distribution as a score⁷ to measure the quality of the generated samples: $\text{NLL}(\mathbf{x}, \mathbf{y}) = -\log p(\mathbf{x}, \mathbf{y})$. The lower NLL the better.

We compare our method to DP-CGAN and DP-MERF at $(\epsilon, \delta) = (1, 10^{-5})$ in Fig. 2. Many of the generated samples by DP-CGAN fall out of the distribution and some modes are dropped (like the green one in the top right corner). DP-MERF preserves all modes. DP-HP performs better than DP-MERF by placing fewer samples in low density regions as indicated by the low NLL. This is due to the drastic difference in the size of the feature map. DP-MERF used 30,000 random features (i.e., 30,000-dimensional feature map). DP-HP used the 25-th order Hermite polynomials on both sum and product kernel approximation (i.e., $25^2 + 25 = 650$ -dimensional feature map). In this example, as the input is 2-dimensional, it was not necessary to subsample the input dimensions to approximate the outer product kernel.

⁷Note that this is different from the other common measure of computing the negative log-likelihood of the true data given the learned model parameters.

Table 1. α -way marginals evaluated on generated samples with discretized Adult and Census datasets.

<i>Adult</i>	PrivBayes		DP-MERF		DP-HP		<i>Census</i>	PrivBayes		DP-MERF		DP-HP	
	$\epsilon=0.3$	$\epsilon=0.1$	$\epsilon=0.3$	$\epsilon=0.1$	$\epsilon=0.3$	$\epsilon=0.1$		$\epsilon=0.3$	$\epsilon=0.1$	$\epsilon=0.3$	$\epsilon=0.1$	$\epsilon=0.3$	$\epsilon=0.1$
$\alpha=3$	0.446	0.577	0.405	0.480	0.332	0.377	$\alpha=2$	0.180	0.291	0.190	0.222	0.141	0.155
$\alpha=4$	0.547	0.673	0.508	0.590	0.418	0.467	$\alpha=3$	0.323	0.429	0.302	0.337	0.211	0.232

5.2. α -way marginals with discretized tabular data

We compare our method to PrivBayes (Zhang et al., 2017) and DP-MERF. For PrivBayes, we used the published code from (McKenna et al., 2019), which builds on the original code with (Zhang et al., 2018) as a wrapper. We test the model on the discretized Adult and Census datasets. Although these datasets are typically used for classification, we use their inputs only for the task of learning the input distribution. Following (Zhang et al., 2017), we measure α -way marginals of generated samples at varying levels of ϵ -DP with $\delta = 10^{-5}$. We measure the accuracy of each marginal of the generated dataset by the total variation distance between itself and the real data marginal (i.e., half of the L1 distance between the two marginals, when both of them are treated as probability distributions). We use the average accuracy over all marginals as the final error metric for α -way marginals. In Table 1, our method outperforms other two at the stringent privacy regime. See Appendix Sec. G.1 for hyperparameter values we used, and Appendix Sec. G.2 for the impact of γ on the quality of the generated samples. We also show how the selection of D_{prod} affects the accuracy in Appendix Sec. G.5.

5.3. Generalization from synthetic to real data

Following (Chen et al., 2020; Torkzadehmahani et al., 2019; Yoon et al., 2019; Chen et al., 2020; Harder et al., 2021; Cao et al., 2021), we evaluate the quality of the (private and non-private) generated samples from these models using the common approach of measuring performance on downstream tasks. We train 12 different commonly used classifier models using generated samples and then evaluate the classifiers on a test set containing *real* data samples. Each setup is averaged over 5 random seeds. The test accuracy indicates how well the models generalize from the synthetic to the real data distribution and thus, the utility of using private data samples instead of the real ones. Details on the 12 models can be found in Table 10.

Tabular data. First, we explore the performance of DP-HP algorithm on eight different imbalanced tabular datasets with both numerical and categorical input features. The numerical features on those tabular datasets can be either discrete (e.g. age in years) or continuous (e.g. height) and the categorical ones may be binary (e.g. drug vs placebo group) or multi-class (e.g. nationality). The datasets are

described in detail in Appendix Sec. G. As an evaluation metric, we use ROC (area under the receiver characteristics curve) and PRC (area under the precision recall curve) for datasets with binary labels, and F1 score for dataset with multi-class labels. Table 2 shows the average over the 12 classifiers trained on the generated samples (also averaged over 5 independent seeds), where overall DP-HP outperforms the other methods in both the private and non-private settings, followed by DP-MERF.⁸ See Appendix Sec. G.3 for hyperparameter values we used. We also show the non-private MERF and HP results in Table 7 in Appendix.

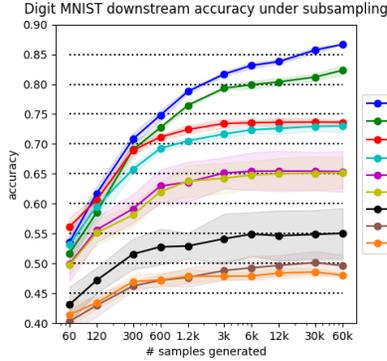
Image data. We follow previous work in testing our method on image datasets MNIST (LeCun et al., 2010) (license: CC BY-SA 3.0) and FashionMNIST (Xiao et al., 2017) (license: MIT). Both datasets contain 60000 images from 10 different balanced classes. We test both fully connected and convolutional generator networks and find that the former works better for MNIST, while the latter model achieves better scores on FashionMNIST. For the experimental setup of DP-HP on the image datasets see Table 9 in Appendix Sec. H.2. A qualitative sample of the generated images for DP-HP and comparison methods is shown in Fig. 4. While qualitatively GS-WGAN produces the cleanest samples, DP-HP outperforms GS-WGAN on downstream tasks. This can be explained by a lack of sample diversity in GS-WGAN shown in Fig. 3.

In Fig. 3, we compare the test accuracy on real image data based on private synthetic samples from DP-GAN, DP-CGAN, GS-WGAN, DP-MERF and DP-HP generators. As additional baselines we include performance of real data and of *full MMD*, a non-private generator, which is trained with the MMD estimator in eq. 2 in a mini-batch fashion. DP-HP gives the best accuracy over the other considered methods followed by DP-MERF but with a considerable difference especially on the MNIST dataset. For GAN-based methods, we use the same weak privacy constraints given in the original papers, because they do not produce meaningful samples at $\epsilon = 1$. Nonetheless, the accuracy these models achieve remains relatively low. Results for individual models for

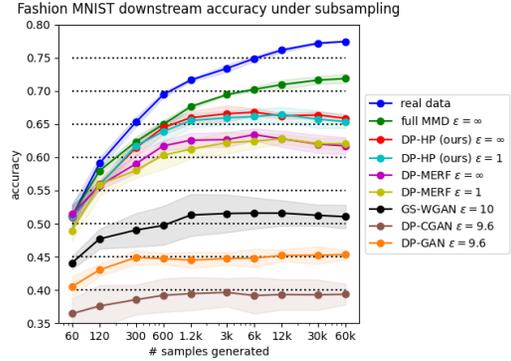
⁸For the Cervical dataset, the non-privately generated samples by DP-MERF and DP-HP give better results than the baseline trained with real data. This may be due to the fact that the dataset is relatively small which can lead to overfitting. The generating samples by DP-MERF and DP-HP could bring a regularizing effect, which improves the performance as a result.

Table 2. Performance comparison on Tabular datasets. The average over five independent runs.

	Real		DP-CGAN ($1, 10^{-5}$)-DP		DP-GAN ($1, 10^{-5}$)-DP		DP-MERF ($1, 10^{-5}$)-DP		DP-HP ($1, 10^{-5}$)-DP	
adult	0.786	0.683	0.509	0.444	0.511	0.445	0.642	0.524	0,688	0,632
census	0.776	0.433	0.655	0.216	0.529	0.166	0.685	0.236	0,699	0,328
cervical	0.959	0.858	0.519	0.200	0.485	0.183	0.531	0.176	0,616	0,312
credit	0.924	0.864	0.664	0.356	0.435	0.150	0.751	0.622	0,786	0,744
epileptic	0.808	0.636	0.578	0.241	0.505	0.196	0.605	0.316	0,609	0,554
isolet	0.895	0.741	0.511	0.198	0.540	0.205	0.557	0.228	0,572	0,498
	F1		F1		F1		F1		F1	
covtype	0.820		0.285		0.492		0.467		0,537	
intrusion	0.971		0.302		0.251		0,892		0.890	



(a) MNIST



(b) FashionMNIST

Figure 3. We compare the real data test accuracy as a function of training set size for models trained on synthetic data from DP-HP and comparison models. Confidence intervals show 1 standard deviation.

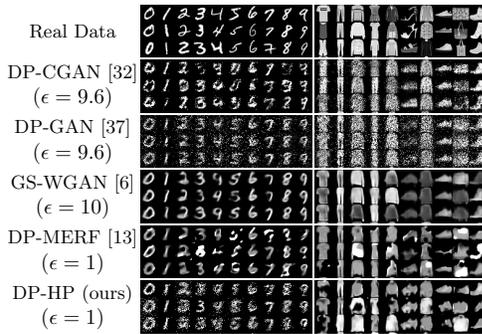


Figure 4. Generated MNIST and FashionMNIST samples from DP-HP and comparison models

both image datasets are given in Appendix Sec. H.

Finally, we show the downstream accuracy for smaller generated datasets down to 60 samples (or 0.1% of original dataset) in Fig. 3. The points, at which additional generated data does not lead to improved performance, gives us a sense of the redundancy present in the generated data. We observe that all generative models except *full MMD* see little increase in performance as we increase the number of synthetic data samples to train the classifiers. This indi-

cates that the *effective dataset size* these methods produce lies only at about 5% (3k) to 10% (6k) of the original data. For DP-GAN and DP-CGAN this effect is even more pronounced, showing little to no gain in accuracy after the first 300 to 600 samples respectively on FashionMNIST.

6. Summary and Discussion

We propose a DP data generation framework that improves the privacy-accuracy trade-off using the Hermite polynomials features thanks to the orderedness of the polynomial features. We chose the combination of outer product and sum kernels computational tractability in handling high-dimensional data. The quality of generated data by our method is significantly higher than that by other state-of-the-art methods, in terms of three different evaluation metrics. In all experiments, we observed that assigning ϵ more to ϵ_1 than ϵ_2 and using the sum kernel’s mean embedding as a main objective together with the outer product kernel’s mean embedding as a constraint (weighted by γ) help improving the performance of DP-HP.

As the size of mean embedding grows exponentially with the input dimension under the outer product kernel, we chose to subsample the input dimensions. However, even with the

subsampling, we needed to be careful not to explode the size of the kernel’s mean embedding, which limits the subsampling dimension to be less than 5, in practice. This gives us a question whether there are better ways to approximate the outer product kernel than random sampling across all input dimensions. We leave this for future work.

Acknowledgements

M.A. Charusaie and F. Harder thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for its support. M. Vinaroz thanks the support by the Gibbs Schule Foundation and the Institutional Strategy of the University of Tübingen (ZUK63). K. Adamczewski is grateful for the support of the Max Planck ETH Center for Learning Systems. M. Park thanks the CIFAR AI Chair fund (at AMII) for its support. M. Park also thanks Wittawat Jitkrittum and Danica J. Sutherland for their valuable time discussing the project. We thank our anonymous reviewers for their constructive feedback.

References

Nist 2018 differential privacy synthetic data challenge. <https://www.nist.gov/ct/pscr/open-innovation-prize-challenges/past-prize-challenges/2018-differential-privacy-synthetic>.

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS ’16*, pp. 308–318, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450341394. doi: 10.1145/2976749.2978318.

Acs, G., Melis, L., Castelluccia, C., and De Cristofaro, E. Differentially private mixture of generative neural networks. *IEEE Transactions on Knowledge and Data Engineering*, 31(6):1109–1121, 2018.

Aronszajn, N. Theory of reproducing kernels. *Trans Am Math Soc*, 68(3):337–404, 1950. ISSN 00029947. URL www.jstor.org/stable/1990404.

Avron, H., Sindhvani, V., Yang, J., and Mahoney, M. W. Quasi-monte carlo feature maps for shift-invariant kernels. *Journal of Machine Learning Research*, 17(120):1–38, 2016. URL <http://jmlr.org/papers/v17/14-538.html>.

Balog, M., Tolstikhin, I., and Schölkopf, B. Differentially private database release via kernel mean embeddings. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80 of *Proceedings of*

Machine Learning Research, pp. 423–431. PMLR, July 2018.

Cao, T., Bie, A., Vahdat, A., Fidler, S., and Kreis, K. Don’t generate me: Training differentially private generative models with sinkhorn divergence. In *NeurIPS*, 2021.

Chen, D., Orekondy, T., and Fritz, M. Gs-wgan: A gradient-sanitized approach for learning differentially private generators. In *Advances in Neural Information Processing Systems 33*, 2020.

Chen, R., Xiao, Q., Zhang, Y., and Xu, J. Differentially private high-dimensional data publication via sampling-based inference. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 129–138, 2015.

Dua, D. and Graff, C. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.

Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., and Naor, M. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques*, volume 4004 of *Lecture Notes in Computer Science*, pp. 486–503. Springer, 2006. doi: 10.1007/11761679_29. URL <https://iacr.org/archive/eurocrypt2006/40040493/40040493.pdf>.

Frigerio, L., de Oliveira, A. S., Gomez, L., and Duverger, P. Differentially private generative adversarial networks for time series, continuous, and discrete open data. In *ICT Systems Security and Privacy Protection - 34th IFIP TC 11 International Conference, SEC 2019, Lisbon, Portugal, June 25-27, 2019, Proceedings*, pp. 151–164, 2019. doi: 10.1007/978-3-030-22312-0_11.

Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *Journal of Machine Learning Research*, 13(Mar):723–773, 2012.

Harder, F., Adamczewski, K., and Park, M. DP-MERF: Differentially private mean embeddings with random features for practical privacy-preserving data generation. In Banerjee, A. and Fukumizu, K. (eds.), *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 1819–1827. PMLR, 13–15 Apr 2021. URL <http://proceedings.mlr.press/v130/harder21a.html>.

Hardt, M., Ligett, K., and Mcsherry, F. A simple and practical algorithm for differentially private data release. In Pereira, F., Burges, C. J. C., Bottou, L., and Weinberger,

- K. Q. (eds.), *Advances in Neural Information Processing Systems* 25, pp. 2339–2347. Curran Associates, Inc., 2012.
- LeCun, Y., Cortes, C., and Burges, C. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- McKenna, R., Sheldon, D., and Miklau, G. Graphical-model based estimation and inference for differential privacy. *arXiv preprint arXiv:1901.09136*, 2019.
- Mehler, F. G. Ueber die entwicklung einer function von beliebig vielen variablen nach laplaceschen functionen höherer ordnung. 1866.
- Mohammed, N., Chen, R., Fung, B. C., and Yu, P. S. Differentially private data release for data mining. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '11, pp. 493–501, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0813-7. doi: 10.1145/2020408.2020487.
- Papernot, N., Abadi, M., Erlingsson, U., Goodfellow, I., and Talwar, K. Semi-supervised Knowledge Transfer for Deep Learning from Private Training Data. In *Proceedings of the International Conference on Learning Representations (ICLR)*, April 2017. URL <http://arxiv.org/abs/1610.05755>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems* 32, pp. 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- Qardaji, W., Yang, W., and Li, N. Priview: practical differentially private release of marginal contingency tables. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*, pp. 1435–1446, 2014.
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pp. 1177–1184, 2008.
- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005. ISBN 026218253X.
- Rudin, W. *Fourier Analysis on Groups: Interscience Tracts in Pure and Applied Mathematics, No. 12*. Literary Licensing, LLC, 2013.
- Sarpatwar, K., Shanmugam, K., Ganapavarapu, V. S., Jagmohan, A., and Vaculin, R. Differentially private distributed data summarization under covariate shift. In *Advances in Neural Information Processing Systems*, pp. 14432–14442, 2019.
- Slepian, D. On the symmetrized kronecker power of a matrix and extensions of mehler’s formula for hermite polynomials. *SIAM Journal on Mathematical Analysis*, 3 (4):606–616, 1972.
- Smola, A., Gretton, A., Song, L., and Schölkopf, B. A Hilbert space embedding for distributions. In *ALT*, pp. 13–31, 2007.
- Smola, A. J. and Schölkopf, B. *Learning with kernels*, volume 4. Citeseer, 1998.
- Snoke, J. and Slavković, A. pmse mechanism: differentially private synthetic data with maximal distributional similarity. In *International Conference on Privacy in Statistical Databases*, pp. 138–159. Springer, 2018.
- Sriperumbudur, B. K., Fukumizu, K., and Lanckriet, G. R. Universality, characteristic kernels and rkhs embedding of measures. *Journal of Machine Learning Research*, 12 (7), 2011.
- Torkzadehmahani, R., Kairouz, P., and Paten, B. Dp-cgan: Differentially private synthetic data and label generation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019.
- Unser, M. and Tafti, P. D. *An introduction to sparse stochastic processes*. Cambridge University Press, 2014.
- Wang, Y.-X., Balle, B., and Kasiviswanathan, S. P. Subsampled rényi differential privacy and analytical moments accountant. PMLR, 2019.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- Xiao, Y., Xiong, L., and Yuan, C. Differentially private data release through multidimensional partitioning. In Jonker, W. and Petković, M. (eds.), *Secure Data Management*, pp. 150–168, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. ISBN 978-3-642-15546-8.
- Xie, L., Lin, K., Wang, S., Wang, F., and Zhou, J. Differentially private generative adversarial network. *CoRR*, abs/1802.06739, 2018.

- Yoon, J., Jordon, J., and van der Schaar, M. PATE-GAN: Generating synthetic data with differential privacy guarantees. In *International Conference on Learning Representations*, 2019.
- Zhang, D., McKenna, R., Kotsogiannis, I., Hay, M., Machanavajhala, A., and Miklau, G. Ektelo: A framework for defining differentially-private computations. *SIGMOD*, 2018.
- Zhang, J., Cormode, G., Procopiuc, C. M., Srivastava, D., and Xiao, X. Privbayes: Private data release via bayesian networks. *ACM Transactions on Database Systems (TODS)*, 42(4):1–41, 2017.
- Zhang, Z., Wang, T., Li, N., Honorio, J., Backes, M., He, S., Chen, J., and Zhang, Y. Privsyn: Differentially private data synthesis. In *30th {USENIX} Security Symposium ({USENIX} Security 21)*, 2021.
- Zhu, H., Williams, C. K., Rohwer, R., and Morciniec, M. Gaussian regression and optimal finite dimensional linear models. 1997.
- Zhu, T., Li, G., Zhou, W., and Yu, P. S. Differentially private data publishing and analysis: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 29(8):1619–1638, August 2017. ISSN 1041-4347. doi: 10.1109/TKDE.2017.2697856.

Appendix

A. Effect of length scale on the kernel approximation

Fig. 5 shows the effect of the kernel length scale on the kernel approximation for both HPs and RFs.

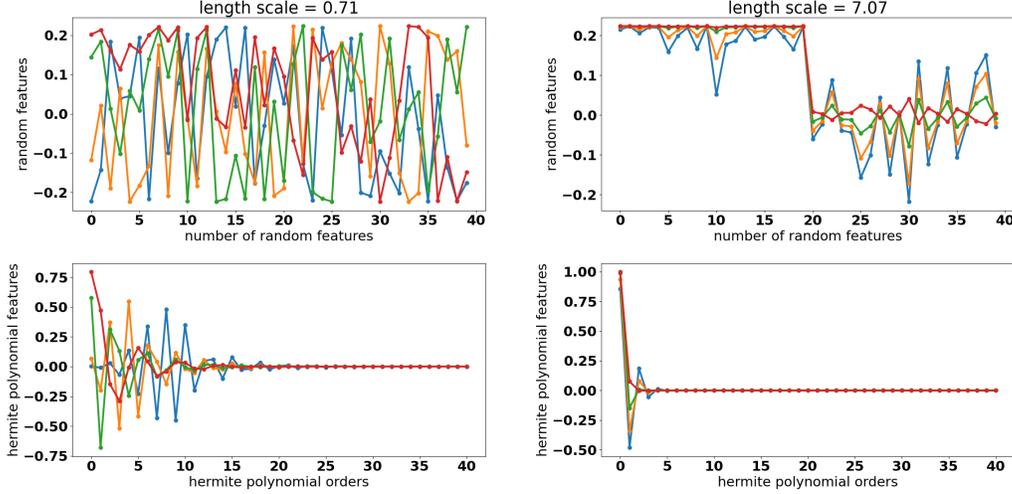


Figure 5. Comparison between HP and random features at a different length scale value. Different color indicates a different datapoint, where four datapoints are drawn from $\mathcal{N}(0, 1)$. **Left:** With length scale $l = 0.71$ (relatively small compared to 1), random features (top) at the four datapoints exhibit large variability while the Hermite polynomial features (bottom) at those datapoints decay at around order ≤ 20 . **Right:** With $l = 7.07$ (large compared to 1), random features (top) exhibit less variability, while it is not clear how many features are necessary to consider. On the other hand, the Hermite polynomial features (bottom) decay fast at around order ≤ 5 and we can make a cut-off at that order without losing much information.

B. Approximation error under HP and Random Fourier features

In the following proposition, we provide that provably our method converges with $O(\rho^{2C})$ where $\rho < 1$ is the constant in the Mehler's formula, while DP-MERF has the convergence $\Omega(1/C)$, where C is the number of features in each case.

Proposition B.1. *Let X and Y be standard normal random variables. There exists a C -dimensional Hermite feature map $\hat{\phi}_{HP}^{(C)}(\cdot)$ with the expected predictive error bounded as*

$$\mathbb{E}_{X,Y} [|k(X, Y) - \langle \hat{\phi}_{HP}^{(C)}(X), \hat{\phi}_{HP}^{(C)}(Y) \rangle|] \leq \frac{1}{3\sqrt{2}} \left(\frac{1}{3}\right)^C. \quad (17)$$

However, the expected predictive error of the random feature map $\hat{\phi}_{RF,\omega}(\cdot)$ with C number of features (i.e., ω is a vector of length C) and the same approximating kernel is equal to

$$\mathbb{E}_{\omega,X,Y} [|k(X, Y) - \langle \hat{\phi}_{RF,\omega}(X), \hat{\phi}_{RF,\omega}(Y) \rangle|] \geq \frac{1}{8C}. \quad (18)$$

Proof. We start by proving eq. 17. In this case, we write the squared error term as following:

$$A_{x,y} = |k(x, y) - \langle \hat{\phi}_{HP}^{(C)}(x), \hat{\phi}_{HP}^{(C)}(y) \rangle|^2 \stackrel{(a)}{=} \left| \sum_{C+1}^{\infty} \frac{\lambda_l}{\sqrt{N_l}} H_l(x) e^{-\frac{\rho}{1+\rho} x^2} \frac{1}{\sqrt{N_l}} H_l(y) e^{-\frac{\rho}{1+\rho} y^2} \right|^2 \quad (19)$$

$$= \sum_{l,l'=C+1}^{\infty} \frac{\lambda_l \lambda_{l'}}{N_l N_{l'}} H_l(x) H_{l'}(x) H_l(y) H_{l'}(y) e^{-\frac{2\rho}{1+\rho} x^2 - \frac{2\rho}{1+\rho} y^2}, \quad (20)$$

where (a) is followed by the definition of $\hat{\phi}_{HP}^{(C)}$ in eq. 6 and its approximation property (i.e., Mehler's formula eq. 5). Now, by setting $\rho = \frac{1}{3}$, we have

$$A_{x,y} = \sum_{l,l'=C+1}^{\infty} \frac{\lambda_l \lambda_{l'}}{N_l N_{l'}} H_l(x) H_{l'}(x) H_l(y) H_{l'}(y) e^{-\frac{1}{2}x^2 - \frac{1}{2}y^2}. \quad (21)$$

Next, we average out $A_{x,y}$ for x s and y s that are drawn from a standard normal distribution as

$$\mathbb{E}_{X,Y \sim N(0,1)} [A_{X,Y}] = \int_{x,y=-\infty}^{\infty} \sum_{l,l'=C+1}^{\infty} \frac{\lambda_l \lambda_{l'}}{N_l N_{l'}} H_l(x) H_{l'}(x) H_l(y) H_{l'}(y) e^{-\frac{1}{2}x^2 - \frac{1}{2}y^2} \frac{e^{-\frac{1}{2}x^2 - \frac{1}{2}y^2}}{2\pi} dx dy \quad (22)$$

$$= \sum_{l,l'=C+1}^{\infty} \frac{\lambda_l \lambda_{l'}}{N_l N_{l'}} \frac{\int H_l(x) H_{l'}(x) e^{-x^2} dx \int H_l(y) H_{l'}(y) e^{-y^2} dy}{2\pi} \quad (23)$$

$$\stackrel{(a)}{=} \sum_{l=C+1}^{\infty} \frac{\lambda_l^2}{N_l^2} \frac{1}{2\pi} \sqrt{\pi} 2^l l! \sqrt{\pi} 2^l l! \stackrel{(b)}{=} \sum_{l=C+1}^{\infty} \frac{(2/3)^2 (1/3)^{2l} 2^{2l} (l!)^2}{\frac{1}{2} 2^{2l} (l!)^2} \frac{2^{2l} (l!)^2}{2} = \frac{4}{9} \sum_{l=C+1}^{\infty} (1/3)^{2l} \quad (24)$$

$$\stackrel{(c)}{=} \frac{1}{2} (1/3)^{2C+2}, \quad (25)$$

where (a) is followed by orthogonality of Hermite polynomials, (b) is followed by the definition of λ_l and N_l in Section 3.1, and (c) is due to the infinite Geometric series.

As a result of eq. 25, the definition of $A_{x,y}$, and Jensen's inequality we have

$$\mathbb{E}_{X,Y} [|k(X,Y) - \langle \hat{\phi}_{HP}^{(C)}(X), \hat{\phi}_{HP}^{(C)}(Y) \rangle|] \leq \mathbb{E}_{X,Y}^{1/2} [A_{X,Y}] \leq \frac{1}{3\sqrt{2}} \left(\frac{1}{3}\right)^C. \quad (26)$$

For bounding the expected error of random features, we expand the squared error using the definition given in eq. 4:

$$B_{x,y,\omega} = \left| k(x,y) - \langle \phi_{RF,\omega}(x), \phi_{RF,\omega}(y) \rangle \right|^2 = \left| e^{-\frac{\rho(x-y)^2}{1-\rho^2}} - \frac{2}{C} \sum_{i=1}^{C/2} \cos \omega_i x \cos \omega_i y - \frac{2}{C} \sum_{i=1}^{C/2} \sin \omega_i x \sin \omega_i y \right| \quad (27)$$

$$= \underbrace{\left| e^{-\frac{\rho(x-y)^2}{1-\rho^2}} - \frac{2}{C} \sum_{i=1}^{C/2} \cos \omega_i (x-y) \right|^2}_{B_{x,y,\omega}}. \quad (28)$$

Next, by setting $\rho = \frac{1}{3}$, we have

$$B_{x,y,\omega} = e^{-\frac{3}{4}(x-y)^2} - \frac{4}{C} e^{-\frac{3}{8}(x-y)^2} \underbrace{\sum_{i=1}^{C/2} \cos \omega_i (x-y)}_{E_{1,x,y,\omega}} + \frac{4}{C^2} \underbrace{\left(\sum_{i=1}^{C/2} \cos \omega_i (x-y) \right)^2}_{E_{2,x,y,\omega}}. \quad (29)$$

Next, we calculate the average of terms $E_{1,x,y,\omega}$ and $E_{2,x,y,\omega}$ over ω .

Due to the Bochner's theorem (see Theorem 3.7 of (Unser & Tafti, 2014)) that shows a shift-invariant positive kernel could be written in the form of Fourier transform of a density function, we have

$$\mathbb{E}_{\omega} [E_{1,x,y,\omega}] = \mathbb{E}_{\omega} \left[\sum_{i=1}^{C/2} \cos \omega_i (x-y) \right] \quad (30)$$

$$= \sum_{i=1}^{C/2} \mathbb{E}_{\omega_i} [e^{j\omega_i(x-y)}] = \frac{C}{2} e^{-\frac{3}{8}(x-y)^2}, \quad (31)$$

Next, we obtain the average of $E_{2,x,y,\omega}$ as following:

$$\mathbb{E}_{\omega} [E_{2,x,y,\omega}] = \mathbb{E}_{\omega} \left[\sum_{i,k=1}^{C/2} \cos \omega_i(x-y) \cos \omega_k(x-y) \right] \quad (32)$$

$$= \mathbb{E}_{\omega} \left[\sum_{i,k=1}^{C/2} \frac{e^{j(\omega_i+\omega_k)(x-y)} + e^{j(\omega_i-\omega_k)(x-y)} + e^{j(-\omega_i+\omega_k)(x-y)} + e^{j(-\omega_i-\omega_k)(x-y)}}{4} \right] \quad (33)$$

$$\stackrel{(a)}{=} \sum_{i,k=1, i \neq k}^{C/2} \mathbb{E}_{\omega_i} [e^{j\omega_i(x-y)}] \mathbb{E}_{\omega_k} [e^{j\omega_k(x-y)}] + \frac{1}{2} \sum_{i=1}^{C/2} (\mathbb{E}_{\omega_i} [e^{j\omega_i(2x-2y)}] + 1) \quad (34)$$

$$\stackrel{(b)}{=} \left(\frac{C^2}{4} - \frac{C}{2} \right) e^{-\frac{3}{4}(x-y)^2} + \frac{C}{4} (e^{-\frac{3}{4}(x-y)^2} + 1) \quad (35)$$

$$= \frac{C^2}{4} e^{-\frac{3}{4}(x-y)^2} + \frac{C}{4} (e^{-\frac{3}{4}(x-y)^2} - 2e^{-\frac{3}{4}(x-y)^2} + 1), \quad (36)$$

where (a) is due to symmetry of the normal distribution of ω , and (b) is followed by independence of ω_i and ω_k and their distribution symmetry.

Substituting eq. 31 and eq. 36 in eq. 29, and using Jensen's inequality, we have

$$\mathbb{E}_{X,Y \sim N(0,1)} \mathbb{E}_{\omega} [B_{x,y,\omega}] = \frac{1}{C} \mathbb{E}_{X,Y \sim N(0,1)} \left[\left(e^{-\frac{3}{4}(X-Y)^2} - 1 \right)^2 \right] \geq \frac{1}{C} \mathbb{E}_{X,Y}^2 \left[e^{-\frac{3}{4}(x-y)^2} - 1 \right] \quad (37)$$

$$= \frac{1}{C} \underbrace{\left(\mathbb{E}_{X,Y \sim N(0,1)} \left[e^{-\frac{3}{4}(X-Y)^2} \right] - 1 \right)^2}_G. \quad (38)$$

To calculate G , we have

$$G = \mathbb{E}_{X,Y \sim N(0,1)} \left[e^{-\frac{3}{4}(X-Y)^2} \right] = \int_{x,y} \frac{e^{-\frac{3}{4}(x^2+y^2-2xy)} e^{-\frac{x^2}{2}-\frac{y^2}{2}}}{2\pi} dx dy \quad (39)$$

$$= \int_{x,y} \frac{e^{-\frac{5}{4}(x^2+y^2)+\frac{3}{2}xy}}{2\pi} dx dy \quad (40)$$

$$= \int_{x,y} \frac{e^{-\frac{5}{4}(x^2-\frac{6}{5}xy+\frac{9}{25}y^2)+\frac{9}{25}\frac{5}{4}y^2-\frac{5}{4}y^2}}{2\pi} dx dy \quad (41)$$

$$\stackrel{(a)}{=} \int_y \frac{e^{-\frac{4}{5}y^2}}{\sqrt{2\pi\frac{5}{2}}} \int_x \frac{e^{-\frac{5}{4}(x-\frac{3}{5}y)^2}}{\sqrt{2\pi\frac{2}{5}}} dx dy \quad (42)$$

$$= \int_y \frac{e^{-\frac{4}{5}y^2}}{\sqrt{2\pi\frac{5}{2}}} dy = \frac{1}{2} \int_y \frac{e^{-\frac{4}{5}y^2}}{\sqrt{2\pi\frac{5}{8}}} dy \quad (43)$$

$$\stackrel{(b)}{=} \frac{1}{2}, \quad (44)$$

where (a) and (b) hold since for a normal distribution $f_{a,b}(x) = \frac{e^{-\frac{(x-b)^2}{2a}}}{\sqrt{2\pi a}}$, we have $\int_x f_{a,b}(x) dx = 1$. As a result of eq. 38 and eq. 44 we have

$$\mathbb{E}_{X,Y,\omega} [B_{X,Y,\omega}] \geq \frac{1}{4C}. \quad (45)$$

Finally, since $0 \leq B_{x,y,\omega} \leq 4$, we have

$$\frac{1}{16C} \leq \mathbb{E}_{X,Y,\omega} \left[\frac{B_{X,Y,\omega}}{4} \right] \leq \mathbb{E}_{X,Y,\omega} \left[\frac{|B_{X,Y,\omega}|^{1/2}}{2} \right] = \frac{1}{2} \mathbb{E}_{X,Y,\omega} [|k(X,Y) - \langle \phi_{RF,\omega}(X), \phi_{RF,\omega}(Y) \rangle |], \quad (46)$$

which proves eq. 18. \square

C. Mercer's theorem and the generalized Hermite polynomials

We first review Mercer's theorem, which is a fundamental theorem on how can we find the approximation of a kernel via finite-dimensional feature maps.

Theorem C.1 ((Smola & Schölkopf, 1998) Theorem 2.10 and Proposition 2.11). *Suppose $k \in L_\infty(\mathcal{X}^2)$, is a symmetric real-valued function, for a non-empty set \mathcal{X} , such that the integral operator $T_k f(x) = \int_{\mathcal{X}} k(x, x') f(x') d\mu(x')$ is positive definite. Let $\psi_j \in L_2(\mathcal{X})$ be the normalized orthogonal eigenfunctions of T_k associated with the eigenvalues $\lambda_j > 0$, sorted in non-increasing order, then*

1. $(\lambda_j)_j \in \ell_1$,
2. $k(x, x') = \sum_{j=1}^{N_{\mathcal{H}}} \lambda_j \psi_j(x) \psi_j(x')$ holds for almost all (x, x') . Either $N_{\mathcal{H}} \in \mathbb{N}$, or $N_{\mathcal{H}} = \infty$; in the latter case, the series converge absolutely and uniformly for almost all (x, x') .

Furthermore, for every $\epsilon > 0$, there exists n such that

$$|k(x, x') - \sum_{j=1}^n \lambda_j \psi_j(x) \psi_j(x')| < \epsilon, \quad (47)$$

for almost all $x, x' \in \mathcal{X}$.

This theorem states that one can define a feature map

$$\Phi_n(x) = [\sqrt{\lambda_1} \psi_1(x), \dots, \sqrt{\lambda_n} \psi_n(x)]^T \quad (48)$$

such that the Euclidean inner product $\langle \Phi(x), \Phi(x') \rangle$ approximates $k(x, x')$ up to an arbitrarily small factor ϵ .

By means of uniform convergence in Mercer's theorem, we can prove the convergence of the approximated MMD using the following lemma.

Lemma C.1. *Let \mathcal{H} be an RKHS that is generated by the kernel $k(\cdot, \cdot)$, and let $\widehat{\mathcal{H}}_n$ be an RKHS with a kernel $k_n(\mathbf{x}, \mathbf{y})$ that can uniformly approximate $k(\mathbf{x}, \mathbf{y})$. Then, for a positive real value ϵ , there exists n , such that for every pair of distributions P, Q , we have*

$$|\text{MMD}_{\mathcal{H}}^2(P, Q) - \text{MMD}_{\widehat{\mathcal{H}}_n}^2(P, Q)| < \epsilon. \quad (49)$$

Proof. Firstly, using Theorem C.1, we can find n such that $|k(x, y) - \langle \Phi_n(x), \Phi_n(y) \rangle| < \frac{\epsilon}{4}$. We define the RKHS $\widehat{\mathcal{H}}_n$ as the space of functions spanned by $\Phi_n(\cdot)$. Next, we rewrite $\text{MMD}_{\mathcal{H}}^2(P, Q) - \text{MMD}_{\widehat{\mathcal{H}}_n}^2(P, Q)$, using the definition of MMD in Section 2.1, as

$$\begin{aligned} & \text{MMD}_{\mathcal{H}}^2(P, Q) - \text{MMD}_{\widehat{\mathcal{H}}_n}^2(P, Q) \\ &= \mathbb{E}_{x, x' \sim P} [k(x, x')] + \mathbb{E}_{y, y' \sim Q} [k(y, y')] - 2\mathbb{E}_{x \sim P, y \sim Q} [k(x, y)] \\ & \quad - \mathbb{E}_{x, x' \sim P} [\langle \Phi_n(x), \Phi_n(x') \rangle] + \mathbb{E}_{y, y' \sim Q} [\langle \Phi_n(y), \Phi_n(y') \rangle] - 2\mathbb{E}_{x \sim P, y \sim Q} [\langle \Phi_n(x), \Phi_n(y) \rangle] \end{aligned} \quad (50)$$

Therefore, we can bound $|\text{MMD}_{\mathcal{H}}^2(P, Q) - \text{MMD}_{\widehat{\mathcal{H}}_n}^2(P, Q)|$ as

$$\begin{aligned} & |\text{MMD}_{\mathcal{H}}^2(P, Q) - \text{MMD}_{\widehat{\mathcal{H}}_n}^2(P, Q)| \stackrel{(a)}{\leq} \left| \mathbb{E}_{x, x' \sim P} [k(x, x')] - \mathbb{E}_{x, x' \sim P} [\langle \Phi_n(x), \Phi_n(x') \rangle] \right| \\ & \quad + \left| \mathbb{E}_{y, y' \sim Q} [k(y, y')] - \mathbb{E}_{y, y' \sim Q} [\langle \Phi_n(y), \Phi_n(y') \rangle] \right| + 2 \left| \mathbb{E}_{x, y \sim P, Q} [k(x, y)] - \mathbb{E}_{x, y \sim P, Q} [\langle \Phi_n(x), \Phi_n(y) \rangle] \right| \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(b)}{\leq} \mathbb{E}_{x,x' \sim P} \left[\left| k(x, x') - \langle \Phi_n(x), \Phi_n(x') \rangle \right| \right] + \mathbb{E}_{y,y' \sim Q} \left[\left| k(y, y') - \langle \Phi_n(y), \Phi_n(y') \rangle \right| \right] \\
 &\quad + 2\mathbb{E}_{x,y \sim P,Q} \left[\left| k(x, y) - \langle \Phi_n(x), \Phi_n(y) \rangle \right| \right] \\
 &\stackrel{(c)}{\leq} \mathbb{E}_{x,x' \sim P} \left[\frac{\epsilon}{4} \right] + \mathbb{E}_{y,y' \sim Q} \left[\frac{\epsilon}{4} \right] + 2\mathbb{E}_{x,y \sim P,Q} \left[\frac{\epsilon}{4} \right] = \epsilon
 \end{aligned} \tag{51}$$

where (a) holds because of triangle inequality, (b) is followed by Tonelli's theorem and Jensen's inequality for absolute value function, and (c) is correct because of the choice of n as mentioned earlier in the proof. \square

As a result of the above theorems, we can approximate the MMD in RKHS \mathcal{H}_k for a kernel $k(\cdot, \cdot)$ via MMD in RKHS $\widehat{\mathcal{H}}_n \subseteq \mathbb{R}^n$ that is spanned by the first n eigenfunctions weighted by square roots of eigenvalues of the kernel $k(\cdot, \cdot)$. Therefore, in the following section, we focus on finding the eigenfunctions/eigenvalues of a multivariate Gaussian kernel.

C.1. Generalized Mehler's approximation

As we have already seen in eq. 5, Mehler's theorem provides us with an approximation of a one-dimensional Gaussian kernel in terms of Hermite polynomials. To generalize Mehler's theorem to a uniform convergence regime (that enables us to approximate MMD via such feature maps as shown in Lemma C.1), and for a multivariate Gaussian kernel, we make use of the following theorem.

Theorem C.2 ((Slepian, 1972), Section 6). *Let the joint Gaussian density kernel $k(\mathbf{x}, \mathbf{y}, C) : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ be*

$$k(\mathbf{x}, \mathbf{y}, C) = \frac{1}{(2\pi)^n |C|^{1/2}} \exp\left(-\frac{1}{2}[\mathbf{x}, -\mathbf{y}]C^{-1}[\mathbf{x}, -\mathbf{y}]^T\right) \tag{52}$$

where C is a positive-definite matrix as

$$C = \begin{bmatrix} C_{11} & C_{12} \\ C_{12}^T & C_{22} \end{bmatrix}, \tag{53}$$

in which $C_{ij} \in \mathbb{R}^{n \times n}$ for $i, j \in \{1, 2\}$, and $C_{11} = C_{22}$. Further, let the integral operator be defined with respect to a measure with density

$$w(\mathbf{x}) = \frac{1}{\int k(\mathbf{x}, \mathbf{y}, C) \partial \mathbf{y}}. \tag{54}$$

Then, the orthonormal eigenfunctions and eigenvalues for such kernel are

$$\psi_{\mathbf{k}}(\mathbf{x}) = \sum_{\mathbf{l}: \|\mathbf{l}\|_1 = \|\mathbf{k}\|_1} (\sigma_{\|\mathbf{k}\|_1}(P)^{-1})_{\mathbf{k}\mathbf{l}} \frac{\varphi_{\mathbf{l}}(\mathbf{x}; C_{11})}{\sqrt{\prod_{i=1}^n l_i!}}, \tag{55}$$

and

$$\lambda_{\mathbf{k}} = \prod_{i=1}^n \mathbf{e}_i^{\mathbf{k}_i/2}. \tag{56}$$

Here, $\sigma_p(A)$ is symmetrized Kronecker power of a matrix A , defined as

$$(\sigma_{\|\mathbf{k}\|_1}(A))_{\mathbf{k}\mathbf{l}} = \sqrt{\prod_{i=1}^n k_i! l_i!} \sum_{M \in \mathbb{R}^{n \times n}: M \mathbf{1}_n = \mathbf{k}, \mathbf{1}_n^T M = \mathbf{l}} \frac{\prod_{ij} A_{ij}^{M_{ij}}}{\prod_{ij} M_{ij}!}, \tag{57}$$

for two n -dimensional vectors \mathbf{k} and \mathbf{l} with $\|\mathbf{k}\|_1 = \|\mathbf{l}\|_1$, the vector \mathbf{e} (the matrix P) is formed by eigenvalues (eigenvectors) of $C_{11}^{-1} C_{12}$, and $\varphi_{\mathbf{l}}(\mathbf{x}, A)$ is generalized Hermite functions defined as

$$\varphi_1(\mathbf{x}, A) = \frac{1}{(2\pi)^{n/2} |A|^{1/2}} \frac{\partial^{\|\mathbf{1}\|_1}}{\partial x_1^{l_1} \dots \partial x_n^{l_n}} \exp\left(-\frac{1}{2} \mathbf{x}^T A^{-1} \mathbf{x}\right). \quad (58)$$

The above theorem provides us with eigenfunctions/eigenvalues of a joint Gaussian density function. We utilize this theorem to approximate Mahalanobis kernels (i.e., a generalization of Gaussian radial basis kernels where $A = cI_n$) via Hermite polynomials as follow.

Proposition C.3. A Mahalanobis kernel $k(\mathbf{x}, \mathbf{y}, A) : \mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}$ defined as

$$k(\mathbf{x}, \mathbf{y}, A) = \exp\left(-(\mathbf{x} - \mathbf{y})A(\mathbf{x} - \mathbf{y})^T\right)$$

can be uniformly approximated as

$$k(\mathbf{x}, \mathbf{y}, A) \simeq \left\langle \Phi_N\left(\sqrt{\frac{\alpha^2 - 1}{\alpha}} \sqrt{A} \mathbf{x}\right), \Phi_N\left(\sqrt{\frac{\alpha^2 - 1}{\alpha}} \sqrt{A} \mathbf{y}\right) \right\rangle, \quad (59)$$

where $\Phi(\mathbf{x}) \in N^D$ is defined as a tensor product

$$\Phi_N(\mathbf{x}) = \bigotimes_{i=1}^n [\phi_{k_i}(x_i)]_{k_i=1}^N, \quad (60)$$

where

$$\phi_{k_i}(x_i) = \left(\frac{(\alpha^2 - 1)\alpha^{-k_i}}{\alpha^2 k_i!}\right)^{1/4} \exp\left(\frac{-x_i^2}{\alpha + 1}\right) H_{k_i}(x_i) \quad (61)$$

Remark 1. Using Proposition C.3 and Lemma C.1, we can show that the MMD based on the tensor feature map in eq. 60 and between any two distributions approximates the real MMD based on Gaussian kernel with Mahalanobis norm.

Proof of Proposition C.3. Let $C = \begin{bmatrix} \frac{1}{2} I_n & \frac{1}{2\alpha} I_n \\ \frac{1}{2\alpha} I_n & \frac{1}{2} I_n \end{bmatrix}$, or equivalently $C^{-1} = \begin{bmatrix} \frac{2\alpha^2}{\alpha^2 - 1} I_n & -\frac{2\alpha}{\alpha^2 - 1} I_n \\ -\frac{2\alpha}{\alpha^2 - 1} I_n & \frac{2\alpha^2}{\alpha^2 - 1} I_n \end{bmatrix}$, for $\alpha \in [1, \infty)$.

Since C is positive-definite, we can define a Gaussian density kernel as

$$k(\mathbf{x}, \mathbf{y}, C) = \frac{1}{\left(\frac{\pi\sqrt{\alpha^2 - 1}}{2\alpha}\right)^n} \exp\left(-\frac{\alpha^2}{\alpha^2 - 1} \|\mathbf{x}\|^2 - \frac{\alpha^2}{\alpha^2 - 1} \|\mathbf{y}\|^2 + \frac{2\alpha}{\alpha^2 - 1} \mathbf{y} \cdot \mathbf{x}^T\right). \quad (62)$$

Moreover, we can calculate the integration over all values of \mathbf{y} as

$$\int k(\mathbf{x}, \mathbf{y}, C) \partial \mathbf{y} = \int \frac{\exp(-\|\mathbf{x}\|^2)}{\left(\frac{\pi\sqrt{\alpha^2 - 1}}{2\alpha}\right)^n} \exp\left(-\frac{\|\alpha \mathbf{y} - \mathbf{x}\|^2}{(\alpha^2 - 1)}\right) \partial \mathbf{y} = \frac{\exp(-\|\mathbf{x}\|^2)}{(\pi)^{n/2}}. \quad (63)$$

Next, by setting $w(\mathbf{x}) = \frac{1}{\int k(\mathbf{x}, \mathbf{y}, C) \partial \mathbf{y}}$ and using Theorem C.2, we have

$$\int \frac{1}{(\pi \frac{\alpha^2 - 1}{\alpha^2})^{n/2}} \psi_{\mathbf{k}}(\mathbf{x}) \exp\left(-\frac{\|\alpha \mathbf{y} - \mathbf{x}\|^2}{\alpha^2 - 1}\right) \partial \mathbf{x} = \lambda_{\mathbf{k}} \psi_{\mathbf{k}}(\mathbf{y}). \quad (64)$$

Now to find the eigenfunctions of the Gaussian kernel $k'(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{\alpha \|\mathbf{x} - \mathbf{y}\|^2}{(\alpha^2 - 1)}\right)$, we let $\psi'_{\mathbf{k}}(\mathbf{x}) = \psi_{\mathbf{k}}(\mathbf{x}) \exp\left(\frac{\alpha}{\alpha + 1} \|\mathbf{x}\|^2\right)$ and let the weight function be $w'(\mathbf{x}) = (\pi)^{n/2} \exp\left(-\frac{(\alpha - 1)}{\alpha + 1} \|\mathbf{x}\|^2\right)$. As a result of such as-

sumptions, we see that

$$\begin{aligned} & \int \psi'_{\mathbf{k}}(\mathbf{x})k'(\mathbf{x}, \mathbf{y})w'(\mathbf{x})\partial\mathbf{x} \\ &= \int (\pi)^{n/2}\psi_{\mathbf{k}}(\mathbf{x})\exp\left(-\frac{1}{\alpha^2-1}\|\mathbf{x}\|^2 - \frac{\alpha}{\alpha^2-1}\|\mathbf{y}\|^2 + \frac{2\alpha}{\alpha^2-1}\mathbf{x}\cdot\mathbf{y}^T\right)\partial\mathbf{x} \end{aligned} \quad (65)$$

$$= (\pi)^{n/2}\exp\left(\frac{\alpha}{\alpha+1}\|\mathbf{y}\|^2\right)\int\psi_{\mathbf{k}}(\mathbf{x})\exp\left(-\frac{\|\alpha\mathbf{y}-\mathbf{x}\|^2}{\alpha^2-1}\right)\partial\mathbf{x} \quad (66)$$

$$\stackrel{(a)}{=} (\pi)^{n/2}\exp\left(\frac{\alpha}{\alpha+1}\|\mathbf{y}\|^2\right)\sqrt{\lambda_{\mathbf{k}}}\psi_{\mathbf{k}}(\mathbf{y})\left(\frac{\pi(\alpha^2-1)}{\alpha^2}\right)^{n/2} \quad (67)$$

$$\stackrel{(b)}{=} (\pi)^n\left(\frac{\alpha^2-1}{\alpha^2}\right)^{n/2}\lambda_{\mathbf{k}}\psi'_{\mathbf{k}}(\mathbf{y}), \quad (68)$$

where (a) holds because of eq. 64, and (b) is followed by the definition of $\psi'_{\mathbf{k}}(\mathbf{y})$. As a result, $\psi'_{\mathbf{k}}(\mathbf{x})$ is an eigenfunction of the integral operator with kernel $k'(\mathbf{x}, \mathbf{y})$ and with weight function $w'(\mathbf{x})$.

Equation eq. 68 shows that the eigenvalue of $k'(\mathbf{x}, \mathbf{y})$ corresponding to $\psi_{\mathbf{k}}(\mathbf{x})$ is as

$$\lambda'_{\mathbf{k}} = (\pi)^n\left(\frac{\alpha^2-1}{\alpha^2}\right)^{n/2}\lambda_{\mathbf{k}} \quad (69)$$

Now we show that such eigenfunctions are orthonormal. Deploying the idea in eq. 68, for two eigenfunctions $\psi'_{\mathbf{k}}(\cdot)$ and $\psi'_{\mathbf{l}}(\cdot)$ for fixed vectors $\mathbf{k}, \mathbf{l} \in \mathbb{N}^n$, we have

$$\int\psi'_{\mathbf{k}}(\mathbf{y})\psi'_{\mathbf{l}}(\mathbf{y})w'(\mathbf{y})\partial\mathbf{y} \stackrel{(a)}{=} \int\psi_{\mathbf{k}}(\mathbf{y})\psi_{\mathbf{l}}(\mathbf{y})\frac{(\pi)^{n/2}}{\exp(-\|\mathbf{x}\|^2)}\partial\mathbf{y} \stackrel{(b)}{=} \int\psi_{\mathbf{k}}(\mathbf{y})\psi_{\mathbf{l}}(\mathbf{y})w(\mathbf{y}) \stackrel{(c)}{=} \delta[\mathbf{l}-\mathbf{k}], \quad (70)$$

where (a) is followed by the definition of eigenfunctions $\psi'_{\mathbf{k}}(\cdot)$, $\psi'_{\mathbf{l}}(\cdot)$ and the definition of weight function $w'(\mathbf{x})$, (b) is due to the definition of $w(\mathbf{x})$ and eq. 63, and (c) holds because of orthonormality of $\psi_{\mathbf{k}}$ s as a result of Theorem C.2.

Further, in this case we have $C_{11}^{-1}C_{12} = \frac{1}{\alpha}I_n$, or equivalently $P = I_n$ and $\mathbf{e} = \frac{1}{\alpha}\mathbf{1}_n$. Hence, firstly using eq. 56, one can see that

$$\lambda_{\mathbf{k}} = \alpha^{-\|\mathbf{k}\|/2}. \quad (71)$$

Secondly, in finding symmetrized Kronecker power $\sigma_{\|\mathbf{k}\|_1}(P)$ in eq. 57, for non-diagonal matrices M , the term $\prod_{i,j} P_{ij}^{M_{ij}} = 0$. Further, for a diagonal matrix M , we have $M\mathbf{1}_n = \mathbf{1}_n M$. This induces the fact that

$$\sigma_{\|\mathbf{k}\|_1}(P) = \begin{cases} 0 & \mathbf{k} \neq \mathbf{1}, \\ 1 & \mathbf{k} = \mathbf{1} \end{cases}. \quad (72)$$

This shows that

$$\psi_{\mathbf{1}}(\mathbf{x}) = \frac{\varphi_{\mathbf{1}}(\mathbf{x})}{\sqrt{\prod_{i=1}^n l_i!}}. \quad (73)$$

To find the formulation of eigenfunction $\psi_{\mathbf{k}}(\mathbf{x})$, we can rewrite the term $\varphi_{\mathbf{1}}(\mathbf{x}, C_{11})$ in eq. 55 for $C_{11} = \frac{1}{2}I_n$ as

$$\varphi_{\mathbf{1}}(\mathbf{x}, I) = \frac{1}{(\pi)^{n/2}} \frac{\partial^{\|\mathbf{1}\|_1}}{\partial x_1^{l_1} \dots \partial x_n^{l_n}} \exp\left(-\sum_{i=1}^n x_i^2\right). \quad (74)$$

We note that the exponential function can be written as the product of functions that are only dependent on one variable x_i for $i \in [n]$. Hence, we can rephrase eq. 74 as a product of the derivative of each function as

$$\varphi_{\mathbf{1}}(\mathbf{x}, I) = \prod_{i=1}^n \frac{1}{\sqrt{\pi}} \frac{\partial^{l_i}}{\partial x_i^{l_i}} \exp(-x_i^2). \quad (75)$$

As a result of this equation and the definition of Hermite functions in one dimension, we have

$$\varphi_1(\mathbf{x}, I) = \frac{\exp(-\|\mathbf{x}\|^2)}{(\pi)^{n/2}} \prod_{i=1}^n H_{l_i}(x_i) \quad (76)$$

Hence, we can calculate $\psi'_{\mathbf{k}}(\mathbf{x})$ as

$$\psi'_{\mathbf{k}}(\mathbf{x}) = \frac{1}{\sqrt{(\pi)^n \prod_{i=1}^n k_i!}} \exp\left(\frac{-\|\mathbf{x}\|^2}{\alpha + 1}\right) \prod_{i=1}^n H_{k_i}(x_i). \quad (77)$$

Using above discussion, we see that \mathbf{k} -th element $[\Phi_N(\mathbf{x})]_{\mathbf{k}}$ of the tensor $\Phi_N(x)$, which is defined in the proposition statement, is equal to

$$[\Phi_N(\mathbf{x})]_{\mathbf{k}} = \sqrt{\lambda'_{\mathbf{k}}} \psi'_{\mathbf{k}}(\mathbf{x}). \quad (78)$$

This fact and Theorem C.1 concludes that we can uniformly approximate $k'(\mathbf{x}, \mathbf{y})$ as

$$k'(\mathbf{x}, \mathbf{y}) = \langle \Phi_N(\mathbf{x}), \Phi_N(\mathbf{y}) \rangle. \quad (79)$$

Further, for any positive-definite matrix A , since the singular values of $\sqrt{\frac{\alpha^2-1}{\alpha}}\sqrt{A}$ are bounded, one can uniformly approximate $k''(\mathbf{x}, \mathbf{y}) := \exp(-(\mathbf{x} - \mathbf{y})A(\mathbf{x} - \mathbf{y})^T) = k'\left(\sqrt{\frac{\alpha^2-1}{\alpha}}\sqrt{A}\mathbf{x}, \sqrt{\frac{\alpha^2-1}{\alpha}}\sqrt{A}\mathbf{y}\right)$ as

$$k''(\mathbf{x}, \mathbf{y}) \simeq \left\langle \Phi_N\left(\sqrt{\frac{\alpha^2-1}{\alpha}}\sqrt{A}\mathbf{x}\right), \Phi_N\left(\sqrt{\frac{\alpha^2-1}{\alpha}}\sqrt{A}\mathbf{y}\right) \right\rangle \quad (80)$$

□

D. Sum-kernel upper-bound

Instead of using Generalized Hermite mean embedding which takes a huge amount of memory, one could use an upper bound to the joint Gaussian kernel. We use the inequality of arithmetic and geometric means to prove that.

$$k(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{1}{2l^2}(\mathbf{x} - \mathbf{y})^T(\mathbf{x} - \mathbf{y})\right) = \exp\left(-\frac{1}{2l^2} \sum_{d=1}^D (x_d - y_d)^2\right) \quad (81)$$

$$= \prod_{d=1}^D \exp\left(-\frac{1}{2l^2}(x_d - y_d)^2\right) \quad (82)$$

$$\stackrel{(a)}{\leq} \frac{1}{D} \sum_{d=1}^D \exp\left(-\frac{D}{2l^2}(x_d - y_d)^2\right) \quad (83)$$

$$= \frac{1}{D} \sum_{d=1}^D k_{X_d}(x_d, y_d), \quad (84)$$

where (a) holds due to inequality of arithmetic and geometric means (AM-GM), and $k_{X_d}(\cdot, \cdot)$ is defined as

$$k_{X_d}(x_d, y_d) := \exp\left(-\frac{D}{2l^2}(x_d - y_d)^2\right). \quad (85)$$

Next, we approximate such kernel via an inner-product of the feature maps

$$\phi_C(\mathbf{x}) = \begin{bmatrix} \phi_{HP,1}^{(C)}(x_1)/\sqrt{D} \\ \phi_{HP,2}^{(C)}(x_2)/\sqrt{D} \\ \vdots \\ \phi_{HP,D}^{(C)}(x_D)/\sqrt{D} \end{bmatrix} \in \mathbb{R}^{((C+1) \cdot D) \times 1}. \quad (86)$$

Although such feature maps are not designed to catch correlation among dimensions, they provide us with a guarantee on marginal distributions as follows.

Lemma D.1. Define $k_{X_i}(\cdot, \cdot)$ as in eq. 85 and define $\phi_C(\mathbf{x})$ as in eq. 86. For $\epsilon \in \mathbb{R}^+$, there exists N such that for $C \geq N$ we have

- $\|\mathbb{E}_{\mathbf{x} \sim P}[\phi_C(\mathbf{x})] - \mathbb{E}_{\mathbf{y} \sim Q}[\phi_C(\mathbf{y})]\|_2 \leq \epsilon \Rightarrow \text{MMD}_{k_{X_i}}(P_i, Q_i) \leq \sqrt{D+1}\epsilon$ for every $i \in \{1, \dots, D\}$, and
- $\text{MMD}_{k_{X_i}}(P_i, Q_i) \leq \epsilon$ for every $i \in \{1, \dots, D\} \Rightarrow \|\mathbb{E}_{\mathbf{x} \sim P}[\phi_C(\mathbf{x})] - \mathbb{E}_{\mathbf{y} \sim Q}[\phi_C(\mathbf{y})]\| \leq \sqrt{2}\epsilon$,

where P_i and Q_i are marginal probability distributions corresponding to P and Q , respectively.

Proof. Since $\phi_{HP,i}^{(C)}(x_i)$ has the certain form as in Theorem C.1, then Lemma C.1 shows that we can use such feature maps to uniformly approximate the MMD in an RKHS based on the kernel $k_i(x_i, y_i) = \exp(-\frac{1}{2l^2}(x_i - y_i)^2)$. As a result, there exists N such that for $C \geq N$, we have

$$\left| \|\mathbb{E}_{x_i \sim P_i}[\phi_{HP,i}^{(C)}(x_i)] - \mathbb{E}_{y_i \sim Q_i}[\phi_{HP,i}^{(C)}(y_i)]\|_2^2 - \text{MMD}_{k_{X_i}}^2(P_i, Q_i) \right| \leq D\epsilon^2. \quad (87)$$

Now we prove the first part. Knowing

$$\|\mathbb{E}_{\mathbf{x} \sim P}[\phi_C(\mathbf{x})] - \mathbb{E}_{\mathbf{y} \sim Q}[\phi_C(\mathbf{y})]\|_2 \leq \epsilon, \quad (88)$$

and by the definition of $\phi_C(\cdot)$, we deduce that

$$\|\mathbb{E}_{x_i \sim P_i}[\phi_{HP,i}^{(C)}(x_i)] - \mathbb{E}_{y_i \sim Q_i}[\phi_{HP,i}^{(C)}(y_i)]\|_2^2 \leq \epsilon^2. \quad (89)$$

Using this and eq. 87 we can prove the first part.

Inversely, by setting $\text{MMD}_{k_{X_i}}(P_i, Q_i) \leq \epsilon$ and eq. 87, one sees that

$$\|\mathbb{E}_{x_i \sim P_i}[\phi_{HP,i}^{(C)}(x_i)] - \mathbb{E}_{y_i \sim Q_i}[\phi_{HP,i}^{(C)}(y_i)]\|_2 \leq \sqrt{2}\epsilon. \quad (90)$$

This coupled with the definition of Φ_C completes the second part of lemma. \square

E. ϕ Recursion

$$\begin{aligned} \phi_{k+1}(x) &= ((1+\rho)(1-\rho))^{\frac{1}{4}} \frac{\rho^{\frac{k+1}{2}}}{\sqrt{2^{k+1}(k+1)!}} H_{k+1}(x) \exp\left(-\frac{\rho}{\rho+1}x^2\right), \quad \text{by definition} \\ &= ((1+\rho)(1-\rho))^{\frac{1}{4}} \frac{\rho^{\frac{k+1}{2}}}{\sqrt{2^{k+1}(k+1)!}} [2xH_k(x) - 2kH_{k-1}(x)] \exp\left(-\frac{\rho}{\rho+1}x^2\right), \\ &= \frac{\sqrt{\rho}}{\sqrt{2(k+1)}} 2x\phi_k(x) - \frac{\rho}{\sqrt{k(k+1)}} k\phi_{k-1}(x). \end{aligned} \quad (91)$$

F. Sensitivity of mean embeddings (MEs)

F.1. Sensitivity of ME under the sum kernel

Here we derive the sensitivity of the mean embedding corresponding to the sum kernel.

$$S_{\hat{\mu}_P^s} = \max_{\mathcal{D}, \mathcal{D}'} \|\hat{\mu}_P^s(\mathcal{D}) - \hat{\mu}_P^s(\mathcal{D}')\|_F = \max_{\mathcal{D}, \mathcal{D}'} \left\| \frac{1}{m} \sum_{i=1}^m \mathbf{h}_s(\mathbf{x}_i) \mathbf{f}(\mathbf{y}_i)^T - \frac{1}{m} \sum_{i=1}^m \mathbf{h}_s(\mathbf{x}'_i) \mathbf{f}(\mathbf{y}'_i)^T \right\|_F$$

where $\|\cdot\|_F$ represents the Frobenius norm. Since \mathcal{D} and \mathcal{D}' are neighbouring, then $m-1$ of the summands on each side cancel and we are left with the only distinct datapoints, which we denote as (\mathbf{x}, \mathbf{y}) and $(\mathbf{x}', \mathbf{y}')$. We then apply the triangle

inequality and the definition of \mathbf{f} . As \mathbf{y} is a one-hot vector, all but one column of $\mathbf{h}_s(\mathbf{x})\mathbf{f}(\mathbf{y})^\top$ are 0, so we omit them in the next step:

$$\begin{aligned} S_{\mu_P^s} &= \max_{(\mathbf{x}, \mathbf{y}), (\mathbf{x}', \mathbf{y}')} \left\| \frac{1}{m} \mathbf{h}_s(\mathbf{x})\mathbf{f}(\mathbf{y})^T - \frac{1}{m} \mathbf{h}_s(\mathbf{x}')\mathbf{f}(\mathbf{y}')^T \right\|_F \\ &\leq \max_{(\mathbf{x}, \mathbf{y})} \frac{2}{m} \left\| \mathbf{h}_s(\mathbf{x})\mathbf{f}(\mathbf{y})^T \right\|_F = \max_{\mathbf{x}} \frac{2}{m} \|\mathbf{h}_s(\mathbf{x})\|_2. \end{aligned} \quad (92)$$

We recall the definition of the feature map given in eq. 10,

$$\|\mathbf{h}_s(\mathbf{x})\|_2 = \frac{1}{\sqrt{D}} \left(\sum_{d=1}^D \|\phi_{HP,d}^{(C)}(x_d)\|_2^2 \right)^{\frac{1}{2}}. \quad (93)$$

To bound $\|\mathbf{h}_s(\mathbf{x})\|_2$, we first prove that $\|\phi_{HP,d}^{(C)}(x_d)\|_2^2 \leq 1$. Using Mehler's formula (see eq. 5), and by plugging in $y = x_d$, one can show that

$$1 = \exp\left(-\frac{\rho}{1-\rho^2}(x_d - x_d)^2\right) = \sum_{c=0}^{\infty} \lambda_c f_c(x_d)^2. \quad (94)$$

Using this, we rewrite the infinite sum in terms of the C th-order approximation and the rest of summands to show that

$$1 = \sum_{c=0}^{\infty} \lambda_c f_c^2(x_d) \stackrel{(a)}{=} \|\phi_{HP,d}^{(C)}(x_d)\|_2^2 + \sum_{c=C+1}^{\infty} \lambda_c f_c^2(x) \stackrel{(b)}{\geq} \|\phi_{HP,d}^{(C)}(x_d)\|_2^2, \quad (95)$$

where (a) holds because of the definition of $\phi_{HP,d}^{(C)}(x_d)$ in eq. 6: $\|\phi_{HP,d}^{(C)}(x_d)\|_2^2 = \sum_{c=0}^C \lambda_c f_c^2(x_d)$, and (b) holds, because λ_c and $f_c^2(x)$ are non-negative scalars.

Finally, deploying eq. 92, eq. 93, and eq. 95, we bound the sensitivity as

$$S_{\mu_P} \leq \max_{\mathbf{x}} \frac{2}{m} \|\mathbf{h}_s(\mathbf{x})\|_2 \leq \frac{2}{m\sqrt{D}} \sqrt{D} = \frac{2}{m}. \quad (96)$$

F.2. Sensitivity of ME under the product kernel

Similarly, we derive the sensitivity of the mean embedding corresponding to the product kernel.

$$S_{\hat{\mu}_P^p} = \max_{\mathcal{D}, \mathcal{D}'} \|\hat{\mu}_P^p(\mathcal{D}) - \hat{\mu}_P^p(\mathcal{D}')\|_F \leq \max_{\mathbf{x}} \frac{2}{m} \|\mathbf{h}_p(\mathbf{x}^{D_{prod}})\|_2$$

Given the definition in eq. 8, the L2 norm is given by

$$\frac{2}{m} \|\mathbf{h}_p(\mathbf{x}^{D_{prod}})\|_2 = \frac{2}{m} \prod_{d=1}^{D_{prod}} \|\phi_{HP}^{(C)}(x_d)\|_2, \quad (97)$$

$$\leq \frac{2}{m} \quad (98)$$

where the last line is due to eq. 95.

G. Descriptions on the tabular datasets

In this section we give more detailed information about the tabular datasets used in our experiments. Unless otherwise stated, the datasets were obtained from the UCI machine learning repository (Dua & Graff, 2017).

Adult

Adult dataset contains personal attributes like age, gender, education, marital status or race from the different dataset participants and their respective income as the label (binarized by a threshold set to 50K). The dataset is publicly available at the UCI machine learning repository at the following link: <https://archive.ics.uci.edu/ml/datasets/adult>.

Census

The Census dataset is also a public dataset that can be downloaded via the SDGym package⁹. This is a clear example of an imbalanced dataset since only 12382 of the samples are considered positive out of a total of 199523 samples.

Cervical

The Cervical cancer dataset comprises demographic information, habits, and historic medical records of 858 patients and was created with the goal to identify the cervical cancer risk factors. The original dataset can be found at the following link: <https://archive.ics.uci.edu/ml/datasets/Cervical+cancer+%28Risk+Factors%29>.

Covtype

This dataset contains cartographic variables from four wilderness areas located in the Roosevelt National Forest of northern Colorado and the goal is to predict forest cover type from the 7 possible classes. The data is publicly available at <https://archive.ics.uci.edu/ml/datasets/covertime>.

Credit

The Credit Card Fraud Detection dataset contains the categorized information of credit card transactions made by European cardholders during September 2013 and the corresponding label indicating if the transaction was fraudulent or not. The dataset can be found at: <https://www.kaggle.com/mlg-ulb/creditcardfraud>. The original dataset has a total number of 284807 samples where only 492 of them are frauds. In our experiments, we discarded the feature related to the time the transaction was done. The data is released under a Database Contents License (DbCL) v1.0.

Epileptic

The Epileptic Seizure Recognition dataset contains the brain activity measured in terms of the EEG across time. The dataset can be found at <https://archive.ics.uci.edu/ml/datasets/Epileptic+Seizure+Recognition>. The original dataset contains 5 different labels that we binarized into two: seizure (2300 samples) or not seizure (9200 samples).

Intrusion

The dataset was used for The Third International Knowledge Discovery and Data Mining Tools Competition held at the Conference on Knowledge Discovery and Data Mining, 1999, and can be found at <http://kdd.ics.uci.edu/databases/kddcup99/kddcup99.html>. We used the file named "kddcup.data10percent.gz" that contains the 10% of the original dataset. The goal is to distinguish between intrusion/attack and normal connections categorized in 5 different labels.

Isolet

The Isolet dataset contains the features sounds as spectral coefficients, contour features, sonorant features, pre-sonorant features, and post-sonorant features of the different letters on the alphabet as inputs and the corresponding letter as the label. The original dataset can be found at <https://archive.ics.uci.edu/ml/datasets/isolet>. However, in our experiments we considered this dataset as a binary classification task as we only considered the labels as constants or vowels.

Table 3 summarizes the number of samples, labeled classes and type of different inputs (numerical, ordinal or categorical) for each tabular dataset used in our experiments. The content of the table reflects the results after pre-processing or binarizing the corresponding datasets.

G.1. Hyperparameters for discrete tabular datasets

Here we include the hyperparameters used in obtaining the results obtained in Table 1. In the main text we describe the choices of the Hermitian hyperparameters. In the separate section G.2 we present a broader view over the gamma hyperparameter.

⁹SDGym package website: <https://pypi.org/project/sdgym/>

Table 3. Tabular datasets. Size, number of classes and feature types descriptions.

dataset	# samps	# classes	# features
isolet	4366	2	617 num
covtype	406698	7	10 num, 44 cat
epileptic	11500	2	178 num
credit	284807	2	29 num
cervical	753	2	11 num, 24 cat
census	199523	2	7 num, 33 cat
adult	48842	2	6 num, 8 cat
intrusion	394021	5	8 cat, 6 ord, 26 num

Table 4. Hyperparameters for discrete tabular datasets

	privacy	batch rate	order hermite prod	prod dimension	gamma	order hermite
Adult	$\epsilon = 0.3$	0.1	10	5	1	100
	$\epsilon = 0.1$	0.1	5	7	1	100
Census	$\epsilon = 0.3$	0.01	5	7	0.1	100
	$\epsilon = 0.1$	0.01	5	7	0.1	100

G.2. Gamma hyperparameter ablation study

Here we study the impact of gamma γ hyperparameter on the quality of the generated samples. Gamma describes the weight that is given to the product kernel in relation to the sum kernel. We elaborate on the results from the Table 1 which describe α -way marginals evaluated on generated samples with discretized Census dataset. We fix all the hyperparameters and vary gamma. The Table 5 shows the impact of gamma. The k -way results remain indifferent for $\gamma \leq 1$ but deteriorate for $\gamma > 1$. In this experiment, we set $\epsilon_1 = \epsilon_2 = \epsilon/2$. Here, “order hermite prod” means the HP order for the outer product kernel, “prod dimension” means the number of subsampled input dimensions, and “order hermite” means the HP order for the sum kernel.

Table 5. The impact of gamma hyperparameter.

epsilon	batch rate	order hermite prod	prod dimension	gamma	epochs	3-way	4-way
0.3	0.1	10	5	0.001	8	0.474	0.570
0.3	0.1	10	5	0.01	8	0.473	0.570
0.3	0.1	10	5	0.1	8	0.499	0.597
0.3	0.1	10	5	1	8	0.474	0.570
0.3	0.1	10	5	10	8	0.585	0.671
0.3	0.1	10	5	100	8	0.674	0.757
0.3	0.1	10	5	1000	8	0.676	0.761

G.3. Training DP-HP generator

Here we provide the details of the DP-HP training procedure we used on the tabular data experiments. Table 6 shows the Hermite polynomial order, the fraction of dataset used in a batch, the number of epochs and the undersampling rate we used during training for each tabular dataset.

G.4. Non-private results

We also show the non-private MERF and HP results in Table 7.

Table 6. Tabular datasets. Hyperparameter settings for private constraints $\epsilon = 1$ and $\delta = 10^{-5}$.

data name	batch rate	order hermite	prod dimension	order hermite	gamma
adult	0.1	5	5	100	0.1
census	0.5	5	5	100	0.1
cervical	0.5	13	5	20	1
credit	0.5	7	5	20	1
epileptic	0.1	5	7	20	0.1
isolet	0.5	13	5	150	1
covtype	0.01	7	2	10	1
intrusion	0.01	5	5	7	1

Table 7. Performance comparison on Tabular datasets. The average over five independent runs.

	Real		DP-MERF (non-priv)		DP-HP (non-priv)		DP-CGAN ($1, 10^{-5}$)-DP		DP-GAN ($1, 10^{-5}$)-DP		DP-MERF ($1, 10^{-5}$)-DP		DP-HP ($1, 10^{-5}$)-DP	
adult	0.786	0.683	0.642	0.525	0.673	0.621	0.509	0.444	0.511	0.445	0.642	0.524	0.688	0.632
census	0.776	0.433	0.696	0.244	0.707	0.32	0.655	0.216	0.529	0.166	0.685	0.236	0.699	0.328
cervical	0.959	0.858	0.863	0.607	0.823	0.574	0.519	0.200	0.485	0.183	0.531	0.176	0.616	0.312
credit	0.924	0.864	0.902	0.828	0.89	0.863	0.664	0.356	0.435	0.150	0.751	0.622	0.786	0.744
epileptic	0.808	0.636	0.564	0.236	0.602	0.546	0.578	0.241	0.505	0.196	0.605	0.316	0.609	0.554
isolet	0.895	0.741	0.755	0.461	0.789	0.668	0.511	0.198	0.540	0.205	0.557	0.228	0.572	0.498
	F1		F1		F1		F1		F1		F1		F1	
covtype	0.820		0.601		0.580		0.285		0.492		0.467		0.537	
intrusion	0.971		0.884		0.888		0.302		0.251		0.892		0.890	

G.5. The effect of subsampled input dimensions for the product kernel on Adult dataset

Table 8 shows the 3-way (Left) and 4-way (Right) marginals evaluated at different number of dimensions for the product kernel (D_{prod}) where the rest of hyperparameters are fixed. The results show that increasing the number of dimensions in the product kernel improved the result.

Table 8. Trade-off for subsampling dimensions in the product kernel for Adult dataset.

ϵ	D_{prod}			D_{prod}		
	2	5	7	2	5	7
1	0.367	0.34	0.332	0.466	0.434	0.422

H. Image data

H.1. Results by model

In the following we provide a more detailed description of the downstreams models accuracy over the different methods considered in the image datasets.

H.2. MNIST and fashionMNIST hyper-parameter settings

Here we give a detailed hyper-parameter setup and the architectures used for generating synthetic samples via DP-HP for MNIST and FashionMNIST datasets in Table 9. The non-private version of our method does not exhibit a significant accuracy difference between 2, 3 and 4 subsampled dimensions for the product kernel, so we considered product dimension to be 2 for memory savings. Table 10 summarizes the 12 predictive models hyper-parameters setup for the image datasets trained on the generated samples via DP-HP. In this experiment, we optimize this loss $\min_{\theta} \|\hat{\mu}_P^p - \hat{\mu}_{Q_{\theta}}^p\|_2^2 + \gamma \|\hat{\mu}_P^s - \hat{\mu}_{Q_{\theta}}^s\|_2^2$, where γ is multiplied by the sum kernel's loss.

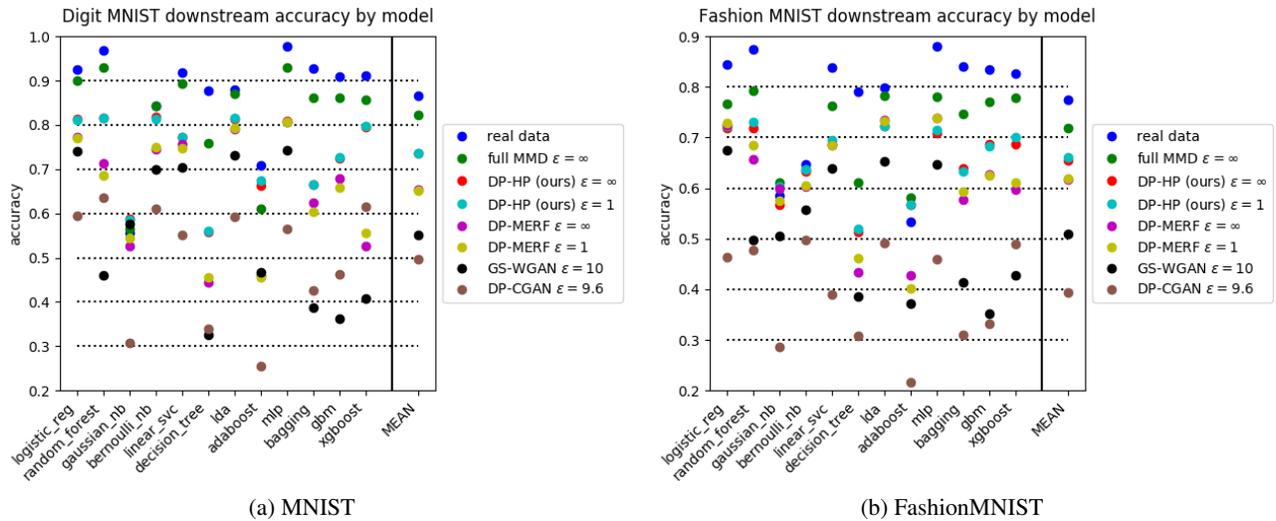


Figure 6. We compare the real data test accuracy of models trained on synthetic data for various models: DP-HP, DP-MERF, GS-WGAN and DP-CGAN. As baselines we also include results for real training data and a generator, which is non-privately trained with MMD, listed as "full MMD". We show accuracy sorted by downstream classifier and the mean accuracy across classifiers on the right. Each score is the average of 5 independent runs.

Table 9. Hyperparameter settings for image data experiments. All parameters not listed here are used with their default values.

	MNIST		FashionMNIST	
	(non-priv)	($1, 10^{-5}$)-DP	(non-priv)	($1, 10^{-5}$)-DP
Hermite order (sum kernel)	100	100	100	100
Hermite order (product kernel)	20	20	20	20
kernel length (sum kernel)	0.005	0.005	0.15	0.15
kernel length (product kernel)	0.005	0.005	0.15	0.15
product dimension	2	2	2	2
subsample product dimension	beginning of each epoch	beginning of each epoch	beginning of each epoch	beginning of each epoch
gamma	5	20	20	10
mini-batch size	200	200	200	200
epochs	10	10	10	10
learning rate	0.01	0.01	0.01	0.01
architecture	fully connected	fully connected	CNN + bilinear upsampling	CNN + bilinear upsampling

Table 10. Hyperparameter settings for downstream models used in image data experiments. Models are taken from the scikit-learn 0.24.2 and xgboost 0.90 python packages and hyperparameters have been set to achieve reasonable accuracies while limiting runtimes. Parameters not listed are kept at their default values.

Model	Parameters
Logistic Regression	solver: lbfgs, max_iter: 5000, multi_class: auto
Gaussian Naive Bayes	-
Bernoulli Naive Bayes	binarize: 0.5
LinearSVC	max_iter: 10000, tol: 1e-8, loss: hinge
Decision Tree	class_weight: balanced
LDA	solver: eigen, n_components: 9, tol: 1e-8, shrinkage: 0.5
AdaBoost	n_estimators: 1000, learning_rate: 0.7, algorithm: SAMME.R
Bagging	max_samples: 0.1, n_estimators: 20
Random Forest	n_estimators: 100, class_weight: balanced
Gradient Boosting	subsample: 0.1, n_estimators: 50
MLP	-
XGB	colsample_bytree: 0.1, objective: multi:softprob, n_estimators: 50