
DP-MERF: Differentially Private Mean Embeddings with Random Features for Practical Privacy-Preserving Data Generation

Frederik Harder^{*1,2}

Kamil Adamczewski^{*1,3}

Mijung Park^{1,2}

¹ Max Planck Institute for Intelligent Systems, Tübingen, Germany

² Department of Computer Science, University of Tübingen, Tübingen, Germany

³ D-ITET, ETH Zurich, Switzerland

{fharder|kadamczewski|mpark}@tue.mpg.de

Abstract

We propose a differentially private data generation paradigm using random feature representations of kernel mean embeddings when comparing the distribution of true data with that of synthetic data. We exploit the random feature representations for two important benefits. First, we require a minimal privacy cost for training deep generative models. This is because unlike kernel-based distance metrics that require computing the kernel matrix on all pairs of true and synthetic data points, we can detach the data-dependent term from the term solely dependent on synthetic data. Hence, we need to perturb the data-dependent term only once and then use it repeatedly during the generator training. Second, we can obtain an analytic sensitivity of the kernel mean embedding as the random features are norm bounded by construction. This removes the necessity of hyper-parameter search for a clipping norm to handle the unknown sensitivity of a generator network. We provide several variants of our algorithm, differentially-private mean embeddings with random features (DP-MERF) to jointly generate labels and input features for datasets such as heterogeneous tabular data and image data. Our algorithm achieves drastically better privacy-utility trade-offs than existing methods when tested on several datasets.

1 Introduction

Differential privacy (DP) is a gold standard privacy notion that is widely used in many applications in machine learning. However, due to its composability, every access to data reduces the privacy guarantee, which limits the number of times one can query sensitive data before a desired privacy level is exceeded. Differentially private data generation solves this problem of limited access by creating a synthetic dataset that is *similar* to the true dataset using DP mechanisms. This process also comes at a privacy cost, but afterwards, the synthetic dataset can be used in place of the true one for unlimited time without further loss of privacy.

Classical approaches to differentially private data generation typically assume a certain class of pre-specified queries. These DP algorithms produce a privacy-preserving synthetic database that is similar to the privacy-sensitive original data for that *fixed query class* [17, 34, 13, 40]. However, specifying a query class in advance, significantly limits the flexibility of the synthetic data, if data analysts hope to perform other machine learning tasks.

To overcome this inflexibility, recent papers on DP data generation have utilized deep generative modelling. The majority of these approaches is based on the generative adversarial networks (GAN) [11] framework, where a discriminator and a generator play a min-max form of game to optimize a given distance metric between the true and synthetic data distributions. Most approaches have used either the *Jensen-Shannon divergence* [20, 30, 36], or the Wasserstein distance [35, 9]. For more details on different divergence metrics, see Supplementary Sec. A.

Another popular choice of distance metric for generative modelling is Maximum Mean Discrepancy (MMD). MMD can compare two probability measures in terms of all possible moments. Therefore, there is no information loss due to a selection of a certain set of moments. The MMD estimator is in closed form (eq. 2) and easy to compute by the pair-wise evaluations of a kernel function using the points drawn from the true and the generated data distributions.

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130. Copyright 2021 by the author(s).

* Equal contribution.

In this work, we propose to use a particular form of MMD via *random Fourier feature* representations [22] of kernel mean embeddings for DP data generation. While MMD can be used within a GAN framework as well (see e.g. [14]) we choose a much simpler method, which is particularly suited for training with DP constraints.

In the objective we use (eq. 3), the mean embedding of the true data distribution (data-dependent) is separate from the embedding of the synthetic data distribution (data-independent). Hence, only the data-dependent term requires privatization. Random features provide an analytic sensitivity of the mean embedding, which allows us to release a DP version of this embedding through a DP mechanism as we explain below. With the privatized data embedding and the synthetic data embedding, our objective no longer directly accesses the data and can be optimized freely to train a data generator. Our contributions are summarized below.

(1) We provide a simple algorithm for DP data generation, which improves on existing methods both in privacy and utility.

- *Simple to optimize:* Since the objective of the optimization contains only a specific private release of data, there are no privacy induced constraints on model choice and optimization method due to privacy. In contrast, methods with private releases as part of the training loop are generally constrained in the number of iterations. As a specific example, DP-SGD requires well-defined sample-wise gradients, which prohibits the use of batch-normalization. Further, increasing the number of trained weights raises the sensitivity of DP-SGD [2] and with it the required strength of gradient perturbation, making large networks infeasible. Our method also avoids the cumbersome min-max optimization present in GAN based approaches and requires only a minimal number of hyperparameters¹.
- *Strong privacy:* Computing the *sensitivity* in our method is *analytically tractable* due to its norm-boundedness of random features. In fact, the norm of random features we use is bounded by 1 by construction. The resulting sensitivity is on the order of 1 over the number of training data points. Consequently, a moderate size of training data can significantly reduce the sensitivity. By requiring only a single DP-release with such a low sensitivity, our method can provide strong DP guarantee more easily than methods which access the data on each training iteration.
- *High utility:* We show in our experiments that our method releases private data with higher utility for downstream tasks than comparison methods. This contrast is particularly stark on MNIST, where our

model at a strong privacy guarantee of $(0.2, 10^{-5})$ -DP outperforms all GAN-based comparison methods, even though they are trained with much weaker privacy of at most $(9.6, 10^{-5})$ -DP.

- *Theoretical study:* We provide an error bound on the objective to theoretically quantify the effect of noise added for privacy to the random feature representation of MMD objective. This bound provides an informative way to select the random feature dimension, given a dataset size and a desired privacy level.

(2) Our algorithm accommodates several needs in privacy-preserving data generation.

- *Generating input and output pairs jointly:* We treat both input and output to be privacy-sensitive. This is different from the conditional-GAN type of methods, where the class distribution is treated as non-sensitive, which increases the risk of successful membership inference, particularly in imbalanced datasets where some classes contain only a small number of samples.
- *Generating imbalanced and heterogeneous tabular data:* Real world datasets may exhibit large variation in data types and class sizes. By addressing both of these issues, we ensure that our algorithm is applicable to a wide variety of datasets.

We start by describing relevant background information in Sec. 2 before introducing our method in Sec. 3 and Sec. 4, followed by an overview of related work in Sec. 5 and experiments in Sec. 6.

2 Background

In the following, we describe the kernel mean embeddings with random features and differential privacy, which our model will use in Sec. 3.

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 [12]

$$\begin{aligned} \text{MMD}^2(P, Q) &= \mathbb{E}_{x, x' \sim P} k(x, x') + \mathbb{E}_{y, y' \sim Q} k(y, y') \\ &\quad - 2\mathbb{E}_{x \sim P} \mathbb{E}_{y \sim Q} k(x, y). \end{aligned} \quad (1)$$

According to the Moore–Aronszajn theorem, there exists a unique Hilbert space \mathcal{H} on which k defines an inner product. Hence, we can find a feature map $\phi: \mathcal{X} \rightarrow \mathcal{H}$ such that $k(x, y) = \langle \phi(x), \phi(y) \rangle_{\mathcal{H}}$, where $\langle \cdot, \cdot \rangle_{\mathcal{H}} = \langle \cdot, \cdot \rangle$ denotes the inner product on \mathcal{H} . Using this fact, we can rewrite the MMD in eq. 1 as [12]

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

where $\mathbb{E}_{x \sim P} [\phi(x)] \in \mathcal{H}$ is known as the (kernel) mean embedding of P , and exists if $\mathbb{E}_{x \sim P} \sqrt{k(x, x)} < \infty$ [25]. The

¹Hyperparameters in our method are the number of random features, a kernel parameter, and the learning rate.

MMD can be interpreted as the distance between the mean embeddings of the two distributions. If k is a characteristic kernel [26], then $P \mapsto \mathbb{E}_{x \sim P}[\phi(x)]$ is injective, and MMD forms a metric, implying that $\text{MMD}(P, Q) = 0$, if and only if $P = Q$.

Given the samples drawn from two probability distributions: $X_m = \{x_i\}_{i=1}^m \sim P$ and $X'_n = \{x'_i\}_{i=1}^n \sim Q$, we can estimate² the MMD by sample averages [12]:

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

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

2.2 Random feature mean embeddings

A fast linear-time MMD estimator can be achieved by considering an approximation to 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 $\hat{\phi}(x) \in \mathbb{R}^D$ and D is the number of features. The resulting approximation of the MMD estimator given in eq. 2 can be computed in $O(m+n)$, i.e., linear in the sample size:

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

One popular approach to obtaining such $\hat{\phi}(\cdot)$ is based on random Fourier features [22] which can be applied to any translation invariant kernel, i.e., $k(x, x') = k(x - x')$ for some function \tilde{k} . According to Bochner's theorem [23], \tilde{k} can be written as $\tilde{k}(x - x') = \int e^{i\omega^\top(x-x')} d\Lambda(\omega) = \mathbb{E}_{\omega \sim \Lambda} \cos(\omega^\top(x - x'))$, where $i = \sqrt{-1}$ and due to positive-definiteness of k , its Fourier transform Λ is non-negative and can be treated as a probability measure. By drawing random frequencies $\{\omega_i\}_{i=1}^D \sim \Lambda$, 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 vector of random Fourier features is given by

$$\hat{\phi}(x) = (\hat{\phi}_1(x), \dots, \hat{\phi}_D(x))^\top \quad (4)$$

where each coordinate is defined by

$$\begin{aligned} \hat{\phi}_j(x) &= \sqrt{2/D} \cos(\omega_j^\top x), \\ \hat{\phi}_{j+D/2}(x) &= \sqrt{2/D} \sin(\omega_j^\top x), \end{aligned}$$

for $j = 1, \dots, D/2$. The approximation error due to these random features was studied in [27].

²Note that this particular MMD estimator is biased.

2.3 Differential privacy

Given privacy parameters $\epsilon \geq 0$ and $\delta \geq 0$, a mechanism \mathcal{M} is (ϵ, δ) -DP if and only if for all possible sets of mechanism outputs S and all neighbouring datasets $\mathcal{D}, \mathcal{D}'$ differing by a single entry, the following equation holds:

$$\Pr[\mathcal{M}(\mathcal{D}) \in S] \leq e^\epsilon \cdot \Pr[\mathcal{M}(\mathcal{D}') \in S] + \delta \quad (5)$$

A DP mechanism guarantees a limit on the amount of information revealed about any one individual in the dataset. Typically this guarantee is achieved by adding randomness to the algorithms' output. Let a function $h : \mathcal{D} \mapsto \mathbb{R}^p$, which is computed on sensitive data \mathcal{D} , output a p -dimensional vector. We can add noise to h for privacy, where the level of noise is calibrated to the *global sensitivity* [8], 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). The *Gaussian mechanism* that we will use in this paper outputs $\tilde{h}(\mathcal{D}) = h(\mathcal{D}) + \mathcal{N}(0, \sigma^2 \Delta_h^2 \mathbf{I}_p)$. The perturbed function $\tilde{h}(\mathcal{D})$ is (ϵ, δ) -DP, where σ is a function of ϵ and δ . For a single application of the mechanism, $\sigma \geq \sqrt{2 \log(1.25/\delta)}/\epsilon$ holds for $\epsilon \leq 1$. The auto-dp package by [31] computes the relationship between ϵ, δ, σ numerically, which we use in our method.

There are two important properties of DP. The *composability* theorem [8] states that the strength of privacy guarantee degrades in a measurable way with repeated use of DP-algorithms. This allows us to combine the results of different private mechanisms in Sec. 4.2 using the advanced composition methods from [32]. Furthermore, the *post-processing invariance* property [8] tells us that the composition of any data-independent mapping with an (ϵ, δ) -DP algorithm is also (ϵ, δ) -DP. This ensures that no analysis of the released synthetic data can yield more information about the real data than what our choice of ϵ and δ allows.

What comes next describes our proposal for privacy-preserving data generation. We first present the vanilla version of our algorithm called, DP-MERF (differentially private mean embeddings with random features).

3 Vanilla DP-MERF for unlabeled data

We first introduce the basic version of our DP-MERF algorithm to learn the distribution of an unlabeled dataset. In this setting, we obtain a data generator by minimizing the random feature representation of MMD, given by

$$\hat{\theta} = \arg \min_{\theta} \widehat{\text{MMD}}_{rf}^2(P_{\mathbf{x}}, Q_{\tilde{\mathbf{x}}_\theta}) \quad (6)$$

where $P_{\mathbf{x}}$ denotes the true data distribution. The samples from Q denoted by $\tilde{\mathbf{x}}$ are drawn from a generative model $\tilde{\mathbf{x}} = G_\theta(\mathbf{z})$. The generative model G_θ is parameterized by θ and takes a sample $\mathbf{z} \sim p(\mathbf{z})$ from a known,

data-independent distribution as input. Using the random Fourier features, we arrive at

$$\widetilde{\text{MMD}}_{r,f}^2(P_{\mathbf{x}}, Q_{\tilde{\mathbf{x}}_\theta}) = \left\| \tilde{\boldsymbol{\mu}}_P - \hat{\boldsymbol{\mu}}_Q \right\|_2^2 \quad (7)$$

where the random feature mean embedding of each distribution is denoted by $\hat{\boldsymbol{\mu}}_P = \frac{1}{m} \sum_{i=1}^m \hat{\phi}(\mathbf{x}_i)$, and $\hat{\boldsymbol{\mu}}_Q = \frac{1}{n} \sum_{i=1}^n \hat{\phi}(G_\theta(\mathbf{z}_i))$.

Notice that $\hat{\boldsymbol{\mu}}_P$ is the only data-dependent term. Hence, we privatize this term by applying the Gaussian mechanism, defining $\tilde{\boldsymbol{\mu}}_P$ by

$$\tilde{\boldsymbol{\mu}}_P = \hat{\boldsymbol{\mu}}_P + \mathcal{N}(0, \Delta_{\tilde{\boldsymbol{\mu}}_P}^2 \sigma^2 I) \quad (8)$$

where the privacy parameter σ is chosen as a function of the privacy budget (ϵ, δ) . The sensitivity of $\hat{\boldsymbol{\mu}}_P$ is analytically tractable due to the triangle inequality and the fact that $\|\hat{\phi}(\cdot)\|_2 = 1$ by construction of the random feature vector given in eq. 4:

$$\Delta_{\tilde{\boldsymbol{\mu}}_P} = \max_{\mathcal{D}, \mathcal{D}'} \left\| \frac{1}{m} \sum_{i=1}^m \hat{\phi}(\mathbf{x}_i) - \frac{1}{m} \sum_{i=1}^m \hat{\phi}(\mathbf{x}'_i) \right\|_2, \quad (9)$$

$$= \max_{\mathbf{x}_n, \mathbf{x}'_n} \left\| \frac{1}{m} \hat{\phi}(\mathbf{x}_n) - \frac{1}{m} \hat{\phi}(\mathbf{x}'_n) \right\|_2 \leq \frac{2}{m}, \quad (10)$$

Due to the post-processing invariance of DP, we can obtain differentially private generator G , since $\hat{\boldsymbol{\mu}}_Q$ is data-independent.

3.1 Bound on the expected absolute error

If we add noise to the random-feature mean embedding of the data distribution, what is the effect of that noise on the learned generator? Theoretically quantifying this effect is challenging under an arbitrary neural network-based generator. Instead, we theoretically quantify the effect of noise on the objective function. In particular, given samples $\mathbf{x} = \{x_i\}_{i=1}^m \sim P$ and $\tilde{\mathbf{x}} = \{\tilde{x}_j\}_{j=1}^n \sim Q$, we want to bound the expected absolute error between the noisy random-feature MMD^2 (eq. 7) and the original estimator MMD^2 (eq. 2). Given the samples, the error deals with two types of randomness. The first arises due to the random features, $\hat{\phi}$. The second arises due to the noise, \mathbf{n} , that we add to the mean-embedding of the data distribution for privacy. The following proposition formally states the bound to the error (See Supplementary Sec. B for proof).

Proposition 3.1. *Given samples $\mathbf{x} = \{x_i\}_{i=1}^m \sim P$ and $\tilde{\mathbf{x}} = \{\tilde{x}_j\}_{j=1}^n \sim Q$, the expected absolute error between the noisy random-feature MMD^2 given in eq. 7 and the MMD^2 given in eq. 2 is bounded by*

$$\mathbb{E}_{\mathbf{n}} \mathbb{E}_{\hat{\phi}} \left[\left| \widetilde{\text{MMD}}_{r,f}^2(\mathbf{x}, \tilde{\mathbf{x}}) - \text{MMD}^2(\mathbf{x}, \tilde{\mathbf{x}}) \right| \right] \quad (11)$$

$$\leq \left(\frac{4D\sigma^2}{m^2} + \frac{8\sqrt{2}\sigma}{m} \frac{\Gamma((D+1)/2)}{\Gamma(D/2)} \right) + 8\sqrt{\frac{2\pi}{D}}. \quad (12)$$

where Γ is the Gamma function, σ is the noise scale (inversely proportional to ϵ), m is the number of training datapoints, and D is the number of features.

Remark 1. *To prove Prop. 3.1, we split eq. 11 into two terms using the triangle inequality. The first term involves the expected absolute error between the **noisy** random feature MMD^2 (eq. 7) and random feature MMD^2 (eq. 3), which yields the first term (inside a big parenthesis) in eq. 12. The second term involves the expected absolute error between random feature MMD^2 (eq. 3) and the MMD^2 (eq. 2), which yields the second term in eq. 12. The upper bound is intuitive in that as the number of random features increases, the second term decreases because the random feature MMD is getting closer to MMD , while the first term increases because we add noise to a larger number of random features.*

Remark 2. *This bound provides a guideline on how to choose D given a desired privacy level ϵ and the dataset size m . First, given m , as long as we choose D such that $m > \sqrt{D}$, the error remains relatively small. However, small D can increase the error in the second term (arising from the MMD approximation using random features). Hence, there is a trade-off between these two terms. In our experiments, the datasets we consider have a relatively large m (see Table 2), and so choosing a large D ($D \approx 10,000$) incurred a relatively small error for a small value of ϵ .*

4 Extension of the vanilla DP-MERF

After introducing the core functionality of DP-MERF, we extend the vanilla method to cases for 1) labeled data, 2) class-imbalanced data, and 3) heterogeneous data.

4.1 DP-MERF for labeled data

We begin by extending our method to balanced labeled datasets with input features \mathbf{x} and output labels \mathbf{y} . In this case, the generator is conditioned on the label: $G_\theta(\mathbf{z}, \mathbf{y}) \mapsto \tilde{\mathbf{x}}$, where \mathbf{y} is drawn from the uniform distribution over classes.

We encode the class information in the MMD objective, by constructing a kernel from a product of two existing kernels, $k((\mathbf{x}, \mathbf{y}), (\mathbf{x}', \mathbf{y}')) = k_{\mathbf{x}}(\mathbf{x}, \mathbf{x}')k_{\mathbf{y}}(\mathbf{y}, \mathbf{y}')$, where $k_{\mathbf{x}}$ is a kernel for input features and $k_{\mathbf{y}}$ is a kernel for output labels. We choose the Gaussian kernel³ for $k_{\mathbf{x}}$ and the polynomial kernel with order-1, $k_{\mathbf{y}}(\mathbf{y}, \mathbf{y}') = \mathbf{y}^\top \mathbf{y}' + c$ for one-hot-encoded labels \mathbf{y} and set $c = 0$. In this case, the result-

³The optimal choice of kernel requires knowledge on the characteristics of the data (see guidelines in Ch. 4 in [33]). At small data sample sizes, a bad kernel choice will affect the efficiency of the algorithm and can underestimate MMD if the chosen kernel assigns small weights to the ‘‘correct’’ frequencies at which the distributions differ. However, with a large enough sample, any characteristic kernel is able to capture such differences.

ing kernel is also characteristic, forming the corresponding MMD as a metric, as explained in [28]. We represent the mean embeddings using random features by

$$\begin{aligned}\hat{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}} &= \frac{1}{m} \sum_{i=1}^m \hat{\mathbf{f}}(\mathbf{x}_i, \mathbf{y}_i), \text{ for true data} \\ \hat{\boldsymbol{\mu}}_{Q_{\mathbf{x},\mathbf{y}}} &= \frac{1}{n} \sum_{i=1}^n \hat{\mathbf{f}}(G_{\theta}(\mathbf{z}_i, \mathbf{y}_i), \mathbf{y}_i), \text{ for synthetic data}\end{aligned}\quad (13)$$

where we define $\hat{\mathbf{f}}(\mathbf{x}_i, \mathbf{y}_i) := \hat{\phi}(\mathbf{x}_i) \mathbf{f}(\mathbf{y}_i)^\top$, where $\mathbf{f}(\mathbf{y}_i) = \mathbf{y}_i$ for the order-1 polynomial kernel and \mathbf{y}_i is one-hot-encoded. See Supplementary Sec. C for derivation. With D random features and C classes, the random feature mean embedding in eq. 13 can also be written as $\hat{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}} = [\mathbf{u}_1, \dots, \mathbf{u}_C] \in \mathbb{R}^{D \times C}$ where c 'th column is given by

$$\mathbf{u}_c = \frac{1}{m} \sum_{\mathbf{x}_i \in X_m^{(c)}} \hat{\phi}(\mathbf{x}_i) \quad (14)$$

where $X_m^{(c)}$ is the set of the datapoints that belong to the class c . As in the unlabeled case, $\hat{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}}$ has sensitivity $\Delta_{\mu_P} = \frac{2}{m}$ and is released with the Gaussian mechanism:

$$\tilde{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}} = \hat{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}} + \mathcal{N}(0, \Delta_{\mu_P}^2 \sigma^2 \mathbf{I}_D) \quad (15)$$

With the released mean embedding $\tilde{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}}$, we construct the private joint maximum mean discrepancy objective:

$$\widetilde{\text{MMD}}_{rf}^2(P_{\mathbf{x},\mathbf{y}}, Q_{\tilde{\mathbf{x}}_{\theta}, \tilde{\mathbf{y}}_{\theta}}) = \left\| \tilde{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}} - \hat{\boldsymbol{\mu}}_{Q_{\mathbf{x},\mathbf{y}}} \right\|_F^2, \quad (16)$$

where F denotes the Frobenius norm. This kind of objective has been used in the non-private setting [39, 10].

4.2 DP-MERF for imbalanced data

Building on the previous section, notice that in eq. 14 the sum in each column is over m_c , the number of instances that belong to the particular class c , while the divisor is the number of samples in the entire dataset, m . This causes difficulties in learning when classes are highly imbalanced, as for rare classes m can be significantly larger than the sum of the corresponding column. In order to address this problem, we release the vector of class counts, $\mathbf{m} = [m_1, \dots, m_C]$ using the Gaussian mechanism:

$$\tilde{\mathbf{m}} = \mathbf{m} + \mathcal{N}(0, \Delta_m^2 \sigma^2 \mathbf{I}_C) \quad (17)$$

As changing a datapoint affects at most two class counts, $\Delta_m = \sqrt{2}$. We then modify the released mean embedding by appropriately weighting the embedding for each class:

$$\tilde{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}}^* = \left[\frac{m}{m_1} \tilde{\mathbf{u}}_1, \dots, \frac{m}{m_C} \tilde{\mathbf{u}}_C \right] \quad (18)$$

Note that we arrive at this expression of mean embedding if we change the kernel on the labels to a weighted one, i.e.,

Algorithm 1 DP-MERF for imbalanced data

Require: Dataset \mathcal{D} , and a privacy level (ϵ, δ)

Ensure: (ϵ, δ) -DP input output samples for all classes

Step 1. Given (ϵ, δ) , compute the privacy parameter σ by the RDP composition in [32] for the two uses of the Gaussian mechanism in steps 2 and 3.

Step 2. Release the mean embedding $\tilde{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}}$ via eq. 15

Step 3. Release the class counts $\tilde{\mathbf{m}}$ using eq. 17.

Step 4. Create the weighted mean embedding $\tilde{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}}^*$ using eq. 18

Step 5. Train the generator by minimizing

$$\widetilde{\text{MMD}}_{rf}^2(P_{\mathbf{x},\mathbf{y}}, Q_{\tilde{\mathbf{x}}_{\theta}, \tilde{\mathbf{y}}_{\theta}}) = \left\| \tilde{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}}^* - \hat{\boldsymbol{\mu}}_{Q_{\mathbf{x},\mathbf{y}}} \right\|_F^2$$

$k_{\mathbf{y}}(\mathbf{y}, \mathbf{y}') = \sum_{c=1}^C \frac{m}{m_c} \mathbf{y}_c^\top \mathbf{y}'_c$. In the re-weighted mean embedding each class-wise embedding $\frac{m}{m_c} \tilde{\mathbf{u}}_c$ has a similar norm, and equally contributes to the objective loss. This ensures that infrequent classes are also modelled accurately.

The total privacy loss results from the composition of the two releases of first $\tilde{\mathbf{m}}$ and then $\tilde{\boldsymbol{\mu}}_{P_{\mathbf{x},\mathbf{y}}}$. During training, we sample the generated labels $\tilde{\mathbf{y}}$ proportional to the class sizes in $\tilde{\mathbf{m}}$. The procedure is summarized in Algorithm 1.

4.3 DP-MERF for heterogeneous data

To handle heterogeneous data consisting of numerical variables denoted by \mathbf{x}_{num} and categorical variables denoted by \mathbf{x}_{cat} , we consider the sum of two existing kernels, $k((\mathbf{x}_{num}, \mathbf{x}_{cat}), (\mathbf{x}'_{num}, \mathbf{x}'_{cat})) = k_{num}(\mathbf{x}_{num}, \mathbf{x}'_{num}) + k_{cat}(\mathbf{x}_{cat}, \mathbf{x}'_{cat})$, where k_{num} is a kernel for numerical variables and k_{cat} is a kernel for categorical variables. Note that this construction of sum of two kernels does not mean that we implicitly assume independence of the two types of variables, for details see Supplementary Sec. I.

As before, we could use the Gaussian kernel for $k_{num}(\mathbf{x}_{num}, \mathbf{x}'_{num}) = \hat{\phi}(\mathbf{x}_{num})^\top \hat{\phi}(\mathbf{x}'_{num})$ and a normalized polynomial kernel with order-1, $k_{cat}(\mathbf{x}_{cat}, \mathbf{x}'_{cat}) = \frac{1}{d_{cat}} \mathbf{x}_{cat}^\top \mathbf{x}'_{cat}$ for one-hot-encoded values \mathbf{x}_{cat} and the length of \mathbf{x}_{cat} being d_{cat} . This normalization is to match the importance of the two kernels in the resulting mean embeddings. Under these kernels, we define

$$\hat{\boldsymbol{\mu}}_{P_{\mathbf{x}}} = \frac{1}{m} \sum_{i=1}^m \hat{\mathbf{h}}(\mathbf{x}_{num}^{(i)}, \mathbf{x}_{cat}^{(i)}), \quad (19)$$

where we define $\hat{\mathbf{h}}(\mathbf{x}_{num}^{(i)}, \mathbf{x}_{cat}^{(i)}) := \left[\begin{array}{c} \hat{\phi}(\mathbf{x}_{num}^{(i)}) \\ \frac{1}{\sqrt{d_{cat}}} \mathbf{x}_{cat}^{(i)} \end{array} \right]$ based on the definition of kernel k (See Supplementary Sec. D for derivation).

In summary, for generating heterogeneous data, we run Algorithm 1 with three changes:

1. Redefine $\hat{\mathbf{f}}(\mathbf{x}, \mathbf{y})$ in eq. 13 as $\hat{\mathbf{h}}(\mathbf{x}_{num}, \mathbf{x}_{cat}) \mathbf{f}(\mathbf{y})^\top$.




Figure 1: 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 three:** Synthetic data generated by DP-CGANs at different privacy levels. CGAN ($\epsilon = \infty$) performs nearly perfectly. However, at $\epsilon = 1$, some modes are dropped, which is reflected in NLL. **Right:** Synthetic data samples generated by DP-MERF at $\epsilon = 1$. Our method captures all modes accurately at $\epsilon = 1$, which is also reflected in NLL.

2. Redefine \mathbf{u}_c in eq. 14 as $\frac{1}{m} \sum_{i \in X_m^{(c)}} \hat{\mathbf{h}}(\mathbf{x}_i)$.
3. Change the sensitivity of \mathbf{u}_c to $\Delta_{\mathbf{u}_c} = \frac{2\sqrt{2}}{m}$ (see Supplementary Sec. G for proof).

5 Related work

Differentially private data release. The field of DP data release contains several distinct lines of research. As mentioned previously, approaches from a learning theory perspective [17, 34, 13, 40] provide bounds on the utility of the data, but contain either strong assumptions about the types of executed queries or intractable computation, which makes this line of research less relevant to our approach.

Among the query-independent methods, a large body of work on DP data release focuses on discrete or possible to discretize data. This is a relevant sub-problem in which good results can be achieved by releasing carefully selected marginals of feature subsets, as each feature only takes on a finite set of values. Such approaches [38, 21, 6] have been, for instance, been dominant among the winning entries of the NIST 2018 Differential Privacy Synthetic Data Challenge [1], which focused on the task of releasing discrete datasets, utilizing related publicly available data. Although we do not compare to this line of work in the main text, as our method deals with the general setting of DP data release, including continuous data, we show the comparison to [38] in the Supplementary Sec. M.

The recent line of research into GAN-based private data release [35, 30, 9, 36, 5] addresses the same general setting and so we select these models for comparison. GANs are regarded as a promising model for this task because of their great success in non-private generative modelling and thanks to the fact the generator network of a GAN can be trained without direct access to the data. The GAN discriminator must still be trained with privacy constraints. In most cases, this is achieved through gradient perturbation using DP-SGD, with the exception of PATE-GAN [36], which

is based on the Private Aggregation of Teacher Ensembles (PATE) [19]. DP-GAN [35] and PATE-GAN [36] generate unlabeled data and thus must train one model per class to obtain a labeled dataset. DP-CGAN [30] and GS-WGAN [5] generate the input features conditioning on the labels, while they do not learn the distribution over the labels. GS-WGAN improves on the basic DP-SGD by alleviating the need for gradient clipping by adapting the loss function and, like PATE-GAN, employs multiple discriminator networks trained on distinct parts of the dataset to amplify privacy by subsampling. We compare these methods with our approach in Sec. 6.

Random feature kernel methods with differential privacy. Some prior work has employed random feature mean embeddings in the context of differential privacy, but not for the purpose of generative modeling. [4] proposed to use the reduced set method in conjunction with random features for sharing DP mean embeddings. This method performs poorly as the dimension of data grows, which is also noted by the authors (see Supplementary Sec. M for comparison to our method). [24] also used the random feature representations of mean embeddings for the DP distributed data summarization to take into account covariate shifts.

6 Experiments

In this section, we show the robustness of our method on a diverse range of data under strong privacy constraints. On each dataset, we train DP-MERF and comparison methods to obtain a set of private *synthetic* data samples and compare, how well these emulate the original dataset. Due to the space limit, we describe all our experimental details (e.g., architecture choices for generators, chosen number of random features, etc.) in the supplementary material. Our code is available at <https://github.com/ParkLabML/DP-MERF>.

Table 1: Performance comparison on tabular datasets, averaged over five runs. DP-MERF achieves the best scores among private models (bold) on the majority of datasets.

	Real		DP-CGAN ($1, 10^{-5}$)-DP		DP-GAN ($1, 10^{-5}$)-DP		DP-MERF ($1, 10^{-5}$)-DP		DP-MERF non-DP	
	ROC	PRC	ROC	PRC	ROC	PRC	ROC	PRC	ROC	PRC
adult	0.730	0.639	0.509	0.444	0.511	0.445	0.650	0.564	0.653	0.570
census	0.747	0.415	0.655	0.216	0.529	0.166	0.686	0.358	0.692	0.369
cervical	0.786	0.493	0.519	0.200	0.485	0.183	0.545	0.184	0.896	0.737
credit	0.923	0.874	0.664	0.356	0.435	0.150	0.772	0.637	0.898	0.774
epileptic	0.797	0.617	0.578	0.241	0.505	0.196	0.611	0.340	0.616	0.335
isolet	0.893	0.728	0.511	0.198	0.540	0.205	0.547	0.404	0.733	0.424
	F1		F1		F1		F1		F1	
covtype	0.643		0.285		0.492		0.467		0.513	
intrusion	0.959		0.302		0.251		0.850		0.856	

2D Gaussian mixtures. We begin our experiments on a simple synthetic distribution of Gaussian mixtures which is aligned on a 5 by 5 grid and assigned to 5 classes as shown in Fig. 1 (left). The dataset is generated by taking 4000 samples from each Gaussian, reserving 10% for the test set, which yields 90000 training samples from the following distribution:

$$p(\mathbf{x}, \mathbf{y}) = \prod_i \sum_{j \in C_{y_i}} \frac{1}{C} \mathcal{N}(\mathbf{x}_i | \boldsymbol{\mu}_j, \sigma \mathbf{I}_2) \quad (20)$$

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 choose this dataset because knowing the true data distribution allows us to compute the negative log likelihood (NLL) of the samples under the true distribution as a measure of the generated samples’ quality: $\text{NLL}(\mathbf{x}, \mathbf{y}) = -\log p(\mathbf{x}, \mathbf{y})$. Note that this is different from the other common measure of computing the negative log-likelihood of the true data given the learned model parameters.

A high NLL score indicates that many samples lie in low density regions of the data distribution. In cases where models tend to under-fit the data, a lower NLL score can thus be regarded as better. However, a low score does not imply that all modes are covered and may also be the result of low sample variance, although the out-of-distribution samples dominate the score, due to the non-linearity of the log function.

At different levels of privacy, we train DP-CGAN on this dataset and select the models with the fewest dropped modes and secondarily the lowest NLL. We compare this to a DP-MERF model for balanced datasets in Fig. 1. While DP-CGAN in the non-private setting ($\epsilon = \infty$) fits the data well, more samples fall out of the distribution as privacy is increased and some modes (like the green one in the top

Table 2: Tabular datasets. num refers to numerical, cat refers to categorical, and ord refers to ordinal variables

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	22561	2	6 num, 8 cat
intrusion	394021	5	8 cat, 6 ord, 26 num

right corner) are dropped. DP-MERF on the other hand preserves all modes and places few samples in low density regions as indicated by the low NLL score. This NLL score is particularly low and on par with the non-private DP-CGAN model, despite a slightly worse fit, because DP-MERF seems to underestimate variance.

Real world data evaluation. In the following experiments we do not know the true data distribution and thus require a different method to evaluate the quality of privately generated datasets. Following the common approach used in [36, 30, 5], we use the private datasets to train a selection of 12 *predictive models* (see Table 5 in the Supplementary for the models). We then evaluate these trained models on a test set of *real* data, which indicates how well the models generalize from the synthetic to the real data distribution and thus how useful the private data would be if used in place of the real data. Note that hyper-parameters of the 12 models differ because the exact settings used in [36] were not available to us, which means that their scores are not directly comparable to ours. As comparison models, we test DP-CGAN [30], as well as our own implementation of an ensemble of 10 DP-GANs, where each model generates data for each class. Our version of DP-GAN differs from [35] in that it uses standard DP-SGD [2] with gradient clipping rather than weight clipping. We further

include GS-WGAN [5] on image datasets following their original setup. Note that our DP-GAN implementation and GS-WGAN use the analytical moments accountant [31] via the autodp package. DP-CGAN uses the RDP accountant [16] from the tensorflow-privacy package, which is slightly older but still comparable. The results in [36, 35] could not be reproduced as the released code was incomplete.

As comparison metrics, we use ROC (area under the receiver operating characteristics curve) and PRC (area under the precision recall curve) for binary-labeled data. For multiclass-labeled data we report accuracy for balanced and F1 score for imbalanced data. As a baseline, we also show the performance of the models trained with the real training data. All the numbers shown in the tables are averages over 5 independent runs.

Table 3: Test accuracy on image data experiments. DP-MERF at $\epsilon = 0.2$ outperforms other methods by a significant margin. $\delta = 10^{-5}$ in all private settings.

	MNIST	FashionMNIST
Real data	0.87	0.78
DP-CGAN $\epsilon = 9.6$	0.50	0.39
DP-GAN $\epsilon = 9.6$	0.48	0.46
GS-WGAN $\epsilon = 10$	0.53	0.50
DP-MERF $\epsilon = 1$	0.65	0.61
DP-MERF $\epsilon = 0.2$	0.61	0.53

Tabular data. We explore the extensions of DP-MERF for imbalanced and heterogeneous data on a number of real-world tabular datasets. These datasets contain numerical features with both discrete and continuous values as well as categorical features with either two classes (e.g. whether a person smokes or not) or several classes (e.g. country of origin). The output labels are also categorical and we include datasets with both binary and multi-class labels. Table 2 summarizes the datasets. Table 1 shows the average across the 12 predictive models trained by the generated samples from DP-CGAN, DP-GAN and DP-MERF. Results for the individual models can be found in Supplementary Sec. K. Overall, our method achieved higher values on the evaluation metrics compared to other methods at the same privacy level.

As a side note, the reason the non-private MERF on Cervical data outperforms the real data is due to the small size of the dataset, which is prone to overfitting. Hence, the added sample variance in the generated data has a regularizing effect and improves the performance.

Image data Finally, we evaluate our method on the image datasets, MNIST and FashionMNIST, which are common benchmarks used in [30, 35, 5]. We apply DP-MERF for balanced data and include convolutional layers, alternating with bi-linear up-sampling, in the generator network to take advantage of the inherent structure of image data.

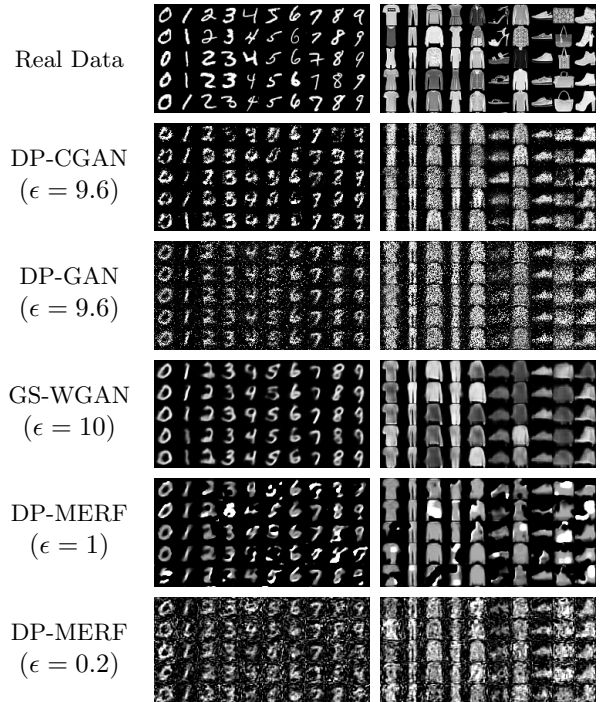


Figure 2: Generated MNIST and FashionMNIST samples from DP-MERF and comparison models with different levels of privacy.

Table 3 compares the test accuracy on real data based on generated samples from DP-CGAN, DP-GAN, GS-WGAN and DP-MERF. Results are averaged over 12 classifiers. For the comparison methods, we use the privacy levels reported in the respective papers, as they do not produce usable samples in the high privacy setting at $\epsilon \leq 1$. It shows that DP-MERF outperforms the GAN based methods by a wide margin and maintains good performance under more meaningful privacy constraints of $(1, 10^{-5})$ -DP and $(0.2, 10^{-5})$ -DP. Low overall scores are largely due to the Adaboost and decision tree models which over-fit to the generated data while other models like logistic regression and multi-layer-perceptrions generalize much better. Detailed results are shown in Supplementary Sec. L.

In the generated samples of the four tested methods in Fig. 2, we see that the samples from DP-MERF at $\epsilon = 0.2$ are noisier than those of GS-WGAN and DP-CGAN, while still achieving higher downstream accuracy.⁴ This indicates that the distinctive features of the data are preserved despite the noisy appearance of the DP-MERF samples. In addition, a loss of sample diversity may explain the worse performance of GS-WGAN and DP-CGAN despite higher perceived sample quality, as we already have observed DP-CGAN dropping modes in the Gaussian data experiment.

⁴As opposed to the version used in [5], the DP-MERF presented here uses an improved generator architecture and privacy analysis, and outperforms GS-WGAN in the classification tasks.

7 Summary and Discussion

We propose a simple and practical algorithm using the random feature representation of kernel mean embeddings for DP data generation. Our method requires a significantly lower privacy budget to produce quality data samples compared to GAN-based approaches, tested on a synthetic dataset, 8 tabular datasets and 2 image datasets. The metrics we use are aimed at supervised learning tasks, but the method is not limited to this application. In the future work, we plan to evaluate our method on a more diverse set of tasks and expand it, to scale to more complex data.

Acknowledgments

We thank Wittawat Jitkrittum, Jia-Jie Zhu, Amin Charusaie and the anonymous reviewers for their valuable time helping us improve our manuscript. All three authors are supported by the Max Planck Society. M. Park and F. Harder are also supported by the Gibbs Schüle Foundation and the Institutional Strategy of the University of Tübingen (ZUK63) and the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039B. F. Harder is grateful for the support of the International Max Planck Research School for Intelligent Systems (IMPRS-IS). K. Adamczewski is grateful for the support of the Max Planck ETH Center for Learning Systems.

References

- [1] Nist 2018 differential privacy synthetic data challenge. <https://www.nist.gov/ctl/pscr/open-innovation-prize-challenges/past-prize-challenges/2018-differential-privacy-synthetic>.
- [2] Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, page 308–318, New York, NY, USA, 2016. Association for Computing Machinery.
- [3] Martín Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *ArXiv*, abs/1701.07875, 2017.
- [4] Matej Balog, Ilya Tolstikhin, and Bernhard Schölkopf. 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*, pages 423–431. PMLR, July 2018.
- [5] Dingfan Chen, Tribhuvanesh Orekondy, and Mario Fritz. Gs-wgan: A gradient-sanitized approach for learning differentially private generators. In *Advances in Neural Information Processing Systems 33*, 2020.
- [6] Rui Chen, Qian Xiao, Yu Zhang, and Jianliang Xu. 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*, pages 129–138, 2015.
- [7] I. Csiszár and P.C. Shields. Information theory and statistics: A tutorial. *Foundations and Trends® in Communications and Information Theory*, 1(4):417–528, 2004.
- [8] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In *Eurocrypt*, volume 4004, pages 486–503. Springer, 2006.
- [9] Lorenzo Frigerio, Anderson Santana de Oliveira, Laurent Gomez, and Patrick Duverger. 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*, pages 151–164, 2019.
- [10] Hongchang Gao and Heng Huang. Joint generative moment-matching network for learning structural latent code. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 2121–2127. International Joint Conferences on Artificial Intelligence Organization, 7 2018.
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.
- [12] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(Mar):723–773, 2012.
- [13] Moritz Hardt, Katrina Ligett, and Frank Mcsherry. A simple and practical algorithm for differentially private data release. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 2339–2347. Curran Associates, Inc., 2012.

- [14] Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabas Poczos. Mmd gan: Towards deeper understanding of moment matching network. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 2203–2213. Curran Associates, Inc., 2017.
- [15] Ryan McKenna, Daniel Sheldon, and Gerome Miklau. Graphical-model based estimation and inference for differential privacy. *arXiv preprint arXiv:1901.09136*, 2019.
- [16] Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275. IEEE, 2017.
- [17] Noman Mohammed, Rui Chen, Benjamin C.M. Fung, and Philip S. Yu. Differentially private data release for data mining. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '11*, pages 493–501, New York, NY, USA, 2011. ACM.
- [18] Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. In *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16*, pages 271–279, USA, 2016. Curran Associates Inc.
- [19] Nicolas Papernot, Martín Abadi, Úlfar Erlingsson, Ian Goodfellow, and Kunal Talwar. Semi-supervised Knowledge Transfer for Deep Learning from Private Training Data. In *Proceedings of the International Conference on Learning Representations (ICLR)*, April 2017.
- [20] Noseong Park, Mahmoud Mohammadi, Kshitij Gorde, Sushil Jajodia, Hongkyu Park, and Youngmin Kim. Data synthesis based on generative adversarial networks. *Proc. VLDB Endow.*, 11(10):1071–1083, June 2018.
- [21] Wahbeh Qardaji, Weining Yang, and Ninghui Li. Priview: practical differentially private release of marginal contingency tables. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*, pages 1435–1446, 2014.
- [22] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in neural information processing systems*, pages 1177–1184, 2008.
- [23] Walter Rudin. *Fourier Analysis on Groups: Inter-science Tracts in Pure and Applied Mathematics, No. 12*. Literary Licensing, LLC, 2013.
- [24] Kanthi Sarpatwar, Karthikeyan Shanmugam, Venkata Sitaramagiridharganesh Ganapavarapu, Ashish Jagmohan, and Roman Vaculin. Differentially private distributed data summarization under covariate shift. In *Advances in Neural Information Processing Systems*, pages 14432–14442, 2019.
- [25] A. Smola, A. Gretton, L. Song, and B. Schölkopf. A Hilbert space embedding for distributions. In *ALT*, pages 13–31, 2007.
- [26] Bharath K Sriperumbudur, Kenji Fukumizu, and Gert RG Lanckriet. Universality, characteristic kernels and rkhs embedding of measures. *Journal of Machine Learning Research*, 12(7), 2011.
- [27] Dougal J. Sutherland and Jeff Schneider. On the error of random fourier features. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence, UAI'15*, page 862–871, Arlington, Virginia, USA, 2015. AUAI Press.
- [28] Zoltán Szabó and Bharath K. Sriperumbudur. Characteristic and universal tensor product kernels. *Journal of Machine Learning Research*, 18(233):1–29, 2018.
- [29] Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Schoelkopf. Wasserstein auto-encoders. In *International Conference on Learning Representations*, 2018.
- [30] Reihaneh Torkzadehmahani, Peter Kairouz, and Benedict Paten. Dp-cgan: Differentially private synthetic data and label generation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019.
- [31] Yu-Xiang Wang, Borja Balle, and Shiva Prasad Kasiviswanathan. Subsampled rényi differential privacy and analytical moments accountant. PMLR, 2019.
- [32] Yu-Xiang Wang, Borja Balle, and Shiva Prasad Kasiviswanathan. Subsampled renyi differential privacy and analytical moments accountant. In Kamalika Chaudhuri and Masashi Sugiyama, editors, *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 1226–1235. PMLR, April 2019.
- [33] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- [34] Yonghui Xiao, Li Xiong, and Chun Yuan. Differentially private data release through multidimensional partitioning. In Willem Jonker and Milan Petković, editors, *Secure Data Management*, pages 150–168, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

- [35] Liyang Xie, Kaixiang Lin, Shu Wang, Fei Wang, and Jiayu Zhou. Differentially private generative adversarial network. *CoRR*, abs/1802.06739, 2018.
- [36] Jinsung Yoon, James Jordon, and Mihaela van der Schaar. PATE-GAN: Generating synthetic data with differential privacy guarantees. In *International Conference on Learning Representations*, 2019.
- [37] Dan Zhang, Ryan McKenna, Ios Kotsogiannis, Michael Hay, Ashwin Machanavajjhala, and Gerome Miklau. Ektelo: A framework for defining differentially-private computations. *SIGMOD*, 2018.
- [38] Jun Zhang, Graham Cormode, Cecilia M Procopiuc, Divesh Srivastava, and Xiaokui Xiao. Privbayes: Private data release via bayesian networks. *ACM Transactions on Database Systems (TODS)*, 42(4):1–41, 2017.
- [39] Yi-Ying Zhang, Chao-Min Shen, Hao Feng, Preston Thomas Fletcher, and Gui-Xu Zhang. Generative adversarial networks with joint distribution moment matching. *Journal of the Operations Research Society of China*, 7(4):579–597, December 2019.
- [40] T. Zhu, G. Li, W. Zhou, and P. S. Yu. Differentially private data publishing and analysis: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 29(8):1619–1638, August 2017.