Differentially Private Variational Autoencoders with Term-wise Gradient Aggregation

Tsubasa Takahashi
LINE Corporation
tsubasa.takahashi@linecorp.com

Shun Takagi *
Kyoto University
takagi.shun.45a@st.kyoto-u.ac.jp

Hajime Ono *
SOKENDAI
hono@ism.ac.jp

Tatsuya Komatsu
LINE Corporation
komatsu.tatsuya@linecorp.com

Abstract

This paper studies how to learn variational autoencoders with a variety of divergences under differential privacy constraints. We often build a VAE with an appropriate prior distribution to describe the desired properties of the learned representations and introduce a divergence as a regularization term to close the representations to the prior. Using differentially private SGD (DP-SGD), which randomizes a stochastic gradient by injecting a dedicated noise designed according to the gradient's sensitivity, we can easily build a differentially private model. However, we reveal that attaching several divergences increase the sensitivity from $O(1)$ to $O(B)$ in terms of batch size B. That results in injecting a vast amount of noise that makes it hard to learn. To solve the above issue, we propose term-wise DP-SGD that crafts randomized gradients in two different ways tailored to the compositions of the loss terms. The term-wise DP-SGD keeps the sensitivity at $O(1)$ even when attaching the divergence. We can therefore reduce the amount of noise. In our experiments, we demonstrate that our method works well with two pairs of the prior distribution and the divergence.

1   Introduction

Privacy-preserving data synthesis (PPDS) is a solution to sharing private data by constructing a generative model while preserving privacy. Differential privacy (DP) [13] is a rigorous notation of privacy to release statistics, and is used in broad domains and applications [3, 6, 9, 11, 27, 29, 31]. In recent years, several works have proposed differentially private deep generative models [4, 20, 30, 33].

Deep generative models have significantly improved in the past few years. Variational autoencoder (VAE) [21, 22] is a likelihood-based model to reconstruct training inputs. VAE also enables us to generate random samples from its learned representations. We often build a VAE with an appropriate prior distribution to describe the desired properties of the representations (such as encouraging clustering, sparsity, and disentanglement), and introduce a divergence as a regularization term to close the learned representations to the prior [5, 8, 16, 17, 24]. This paper studies how to learn variational autoencoders with a variety of divergences under differential privacy constraints.

A simple way to build a differentially private VAE is to employ differentially private stochastic gradient descent (DP-SGD) [2] in the learning process of vanilla VAE. The key idea of DP-SGD is that it injects noises to stochastic gradients for giving DP guarantees to the learned parameters.

---

*Work done at LINE Corporation.

The noise scale is designed according to the stochastic gradient's sensitivity, which is the maximal change of the gradient when any single input is modified. To limit the gradient's sensitivity, DP-SGD first decomposes input samples into disjoint smaller groups (i.e., micro-batches). Then, DP-SGD computes a stochastic gradient for each group and clips the norm of the gradient by a constant. On the other hand, a misuse of the gradient aggregations might cause privacy leakages unconsciously.

Our contributions are three-fold. First, we reveal that several divergences might increase the stochastic gradients' sensitivity when attaching them to the loss function. To discover the issues, we address a sensitivity study in the learning process of VAEs based on DP-SGD. The sensitivity is increased from $O(1)$ to $O(B)$ in terms of batch size B by attaching the divergence. Consequently, the sensitivity increase degrades the quality of the learned model since it directly amplifies the amount of noise. If unfortunately, we do not notice the sensitivity increase, we might cause an insufficient differential privacy guarantee.

Second, to solve the above issue, we propose *term-wise DP-SGD* that crafts randomized gradients in two different ways tailored to the compositions of the loss terms. The term-wise DP-SGD keeps the sensitivity at $O(1)$ even when attaching the divergence. We can therefore build a differentially private VAE with a small amount of noise by our proposed method.

Third, based on the term-wise DP-SGD, we present PRIVAE, a general model to learn VAEs with attaching a variety of divergences while satisfying differential privacy. Our experiments demonstrate that our proposed method works well with two pairs of the prior distribution and the divergence.

This paper clarifies how to aggregate gradients in VAEs to satisfy differential privacy while refraining the amount of noise. Although we mainly study differentially private VAEs, these contributions also have significant importance for the other machine learning models to satisfy differential privacy.

1.1 Related Works

Generative models under differential privacy have been studied in a last decade. Traditional approaches are based on capturing probabilistic models, low rank structure, and learning statistical characteristics from original sensitive database [12, 35, 36]. Plausible deniability [10] is an extended privacy metric behind DP for building a generative model.

We have several studies about DP-SGD [26, 25, 34]. Lee et al. [23] demonstrated that DP-SGD can be improved with adaptive step sizes and careful allocation of privacy budgets between iterations. Bagdasaryan et al. [7] revealed that if the original model is unfair, the unfairness becomes worse once DP is applied.

2 Preliminaries

2.1 Differential Privacy

Differential privacy [13, 14, 15] is a rigorous mathematical privacy definition, which quantitatively evaluates the degree of privacy protection when we publish statistical outputs. The definition of differential privacy is as follows:

Definition 1 ((ε, δ)-differential privacy) *A randomized mechanism $\mathcal{M} : \mathcal{D} \to \mathcal{Z}$ satisfies (ε, δ)-differential privacy if, for any two neighboring input $D, D' \in \mathcal{D}$ and any subset of outputs $Z \subseteq \mathcal{Z}$, it holds that*

$$\Pr[\mathcal{M}(D) \in Z] \leq \exp(\varepsilon) \Pr[\mathcal{M}(D') \in Z] + \delta. \tag{1}$$

Practically, we employ a randomized mechanism $\mathcal{M}$ that ensures differential privacy for a function f. The mechanism $\mathcal{M}$ perturbs the output of f to cover f's sensitivity that is the maximum degree of change over any pairs of D and D'.

Definition 2 (Sensitivity) *The sensitivity of f for any two neighboring input $D, D' \in \mathcal{D}$ is*

$$\Delta_f = \sup_{D, D' \in \mathcal{D}} ||f(D) - f(D')||. \tag{2}$$

where $|| \cdot ||$ is a norm function defined on f's output domain.

Based on the sensitivity of f, we design the degree of noise to ensure differential privacy. Laplace mechanism and Gaussian mechanism are well-known as standard approaches.

Let $M_1, \ldots, M_k$ be mechanisms satisfying $(\varepsilon_1, \delta_1)$-, $\ldots$, $(\varepsilon_k, \delta_k)$-differential privacy, respectively. Then, a mechanism sequentially applying $M_1, \ldots, M_k$ satisfies $(\sum_{i \in [k]} \varepsilon_i, \sum_{i \in [k]} \delta_i)$-differential privacy. This fact refers to *composability*[13]. In particular, this composition is called sequential composition.

2.2 DP-SGD

Differentially private stochastic gradient descent (DP-SGD) [2], is a useful optimization technique for learning a model f under differential privacy constraints. The key idea of DP-SGD is that it adds noise to stochastic gradients during training for making differential privacy guarantees on f's parameters θ. To obtain the scale of noise, DP-SGD limits ℓ_2-sensitivity of stochastic gradient $\mathbf{g}$ by clipping its norm. The gradient clipping π_C that limits the sensitivity up to C is denoted as follows:

$$\pi_C(\mathbf{g}) = \mathbf{g} * \min\left(1, \frac{C}{||\mathbf{g}||_2}\right) \tag{3}$$

In the DP-SGD, we compute an empirical loss for each micro-batch that includes only one sample. For each micro-batch, DP-SGD generates its clipped gradient. Based on the clipped gradients, DP-SGD crafts a randomized gradient $\tilde{\mathbf{g}}$ through computing the average over the clipped gradients and adding noise whose scale is defined by C and σ_ε, where σ_ε is noise scaler to satisfy (ε, δ)-DP.

$$\tilde{\mathbf{g}} = \frac{1}{B}\left(\sum_{i \in [B]} \pi_C(\mathbf{g}_i) + \mathcal{N}(0, (\sigma_\varepsilon C)^2 \mathbf{I})\right). \tag{4}$$

At last, DP-SGD takes a step based on the randomized gradient $\tilde{\mathbf{g}}$. Abadi et al. [2] also proposed a moment accountant that maintain privacy loss more precisely than the sequential composition. In the moment accountant, σ_ε has the following relationship against ε and δ (Theorem 1 in [2]).

$$\sigma_\varepsilon \geq c_2 \frac{q\sqrt{T \log(1/\delta)}}{\varepsilon} \tag{5}$$

where q is a sampling probability, T is a number of steps and c_2 is a constant number. To compute the privacy loss through moment accountant, we can utilize a useful tool in Tensorflow privacy [1].

2.3 Variational Autoencoder

Variational autoencoder (VAE) [22] is a model to learn parametric latent variables by maximizing the marginal log-likelihood of the training data points. VAE consists of two parts, inference model $q(z|x)$ for an encoder $g(x; \theta)$, and the likelihood model $p(x|z)$ for a decoder $f(z; \theta)$.

Variational evidence lower bound.   Introduction of an approximate posterior $q_\phi(z|x)$ enable us to construct variational evidence lower bound (ELBO) on log-likelihood $\log p(x)$ as

$$\begin{aligned}\mathcal{L}_{ELBO} &= \log p(x) - D_{KL}(q(z|x)||p(x|z)) \\ &= \mathbb{E}_{q(z|x)}[\log p(x|z)] - D_{KL}(q(z|x)||p(z)) \\ &\leq \log p(x).\end{aligned} \tag{6}$$

To implement encoder and decoder as a neural network, we need to backpropagate through random sampling. However, such backpropagation does not flow through the random samples. To overcome this issue, VAE introduces the reparametrization trick. The trick can be described as $z = \mu + \Sigma\epsilon$ where $\epsilon \sim \mathcal{N}(0, \mathbf{I})$. After constructing VAE, we can generate random samples following the two steps; 1) choose a latent vector $z \sim \mathcal{N}(0, \sigma^2 I)$, and 2) generate $\tilde{x}$ by decoding z. $\tilde{x} = f(z)$.

Attaching a divergence for regularization.   To capture the desired property in the learned representation space of VAEs, we can employ a variety of prior distributions as $p(z)$ and an additional regularization term. We assume an additional regularization term $D(q(z), p(z))$, that is a divergence between $q(z)$ and $p(z)$. The ELBO with the regularization is described as follows:

$$\mathcal{L}_{\text{ELBO}} = \mathbb{E}_{q(z|x)}[\log p(x|z)] - \beta D_{\text{KL}}(q(z|x)||p(z)) - \alpha D(q(z), p(z)) \tag{7}$$

Several $D(q(z), p(z))$ are difficult to be decomposed into micro-batch losses that DP-SGD requires.

3 Sensitivity Analysis

Here we address a sensitivity study in DP-SGD for VAEs with various loss functions to clarify the required nose scale for ensuring differential privacy on the parameters of VAEs.

3.1 Learning VAEs in DP-SGD

Let $batch = \{x_i\}_{i=1}^B$ is a randomly selected samples with sampling probability B/N. We assume the loss function of VAE is formed as the following abstract equation:

$$\mathcal{L} = -\mathcal{L}_{ELBO} = \mathbb{E}_x\left[\phi(x)\right] + \psi(batch) \tag{8}$$

where $\phi(x_i)$ is a function which computes a loss only depend on x_i, and $\psi(batch)$ is a function which computes loss value across all samples in *batch* ($=\{x_1, \ldots, x_B\}$). We call $\phi(x_i)$ *sample-wise term*, and $\psi(batch)$ *batch-wise term*. The loss function (8) is also rewritten as follows:

$$\mathcal{L} = \frac{1}{B}\sum_{i \in [B]} \mathcal{L}_i = \frac{1}{B}\sum_{i \in [B]}\left(\phi(x_i) + \psi(batch)\right) \tag{9}$$

where $\mathcal{L}_i$ is a micro-batch loss. In DP-SGD, the stochastic gradient of $\mathcal{L}_i$ is clipped by C as (3). That means the sensitivity of the gradient is bounded by the constant. At the last step in a batch, we craft a randomized gradient through aggregating the clipped gradients and injecting noise whose scale is $C\sigma_\varepsilon$ to ensure differential privacy. This aggregation has an effort to reduce the variance of the noise. We call the construction of (9) *micro aggregation*.

Based on the above assumptions, we can see the following series of propositions.

Proposition 1 *Assume $\mathcal{L}_i = \phi(x_i)$ and the stochastic gradient of $\mathcal{L}_i$ is clipped by (3) with the constant C, ℓ_2-sensitivity of $\mathcal{L}$ is C.*

Proof *Let $\boldsymbol{g}_i$ be the stochastic gradient of $\mathcal{L}_i = \phi(x_i)$. Since $\phi(x_i)$ is independent from $\phi(x_j)$ of $j \neq i$, changing x_i only modifies its clipped gradient $\pi_C(\boldsymbol{g}_i)$. Thus, the sensitivity is $1 \times C$.*

Proposition 2 *Assume $\mathcal{L}_i = \psi(batch)$ and the stochastic gradient of $\mathcal{L}_i$ is clipped by (3) with the constant C, ℓ_2-sensitivity of $\mathcal{L}$ is BC.*

Proof *Let $\boldsymbol{g}_i$ be the stochastic gradient of $\mathcal{L}_i = \psi(batch)$. While $\psi(batch)$ is shared in all $\mathcal{L}_j$, $\forall j \in [B]$, the change of x_i modifies all $\mathcal{L}_j$. Thus, the sensitivity is $B \times C$.*

Proposition 3 *Assume $\mathcal{L}_i = \phi(x_i) + \psi(batch)$ and the stochastic gradient of $\mathcal{L}_i$ is clipped by (3) with the constant C, ℓ_2-sensitivity of $\mathcal{L}$ is BC.*

Proof *As well as the Proof of Proposition 2, since $\psi(batch)$ is shared in all $\mathcal{L}_j$, $\forall j \in [B]$, the change of x_i modifies all $\mathcal{L}_j$. Thus, the sensitivity is $B \times C$.*

From the above three propositions, we reach the following theorem about the sensitivity for learning differentially private VAEs in the DP-SGD manner.

Theorem 1 *ℓ_2-sensitivity of $\mathcal{L}$ for learning a vanilla VAE is either BC or C.*

Proof *Let rec_i be the reconstruction loss (i.e., negative log-likelihood) of x_i. The loss functions of vanilla VAE can be written as $\mathcal{L}_i = rec_i + KL(q(z|x)||p(z))$. For this formulation, the sensitivity is BC from the proposition 3. Fortunately, the KL term can be decomposed as follows:*

$$KL(q(z|x)||p(z)) = \frac{1}{B}\sum_{i \in [B]}\left(\log q(z|x_i) - \log p(z)\right) \tag{10}$$

Thus we can rewrite the loss as sample-wise form that does not depend on the other samples:

$$\mathcal{L}_i = rec_i + kld_i \quad where \quad kld_i = \log q(z|x_i) - \log p(z) \tag{11}$$

From Proposition 1, the sensitivity when we utilize (11) is C.

Lemma 1 *Let a VAE introduces an additional regularization term and the regularization term cannot be decomposed into micro-batch losses that every micro-batch depends on an only single input. The sensitivity of DP-SGD for learning the VAE with the regularization is BC.*

On the other hand, DP-SGD is applicable not only for micro-batches but also for the overall batch ($\mathcal{L} = \psi(batch)$). When we craft the randomized gradient from the overall batch, the stochastic gradient' sensitivity keeps at $O(1)$. We call this construction *batch aggregation*. By employing the batch aggregation, we can compute a divergence from all samples in the batch without increasing the sensitivity. However, the batch aggregation also injects a large amount of noise because it does not have a factor to reduce the noise that micro-batch organizations have. Therefore, DP-SGD often organizes the micro-batches whose size is one for crafting the randomized gradient.

3.2 Privacy Leakage

As discussed the above study, ill constructions of the randomized gradient that aggregates micro-batch losses like $\mathcal{L}_i = \text{rec}_i + \text{kld}_i + D(q(z), p(z))$ and injects insufficient scale of noise to cover the increased sensitivity fail into differential privacy guarantee that we expected. In this case, unfortunately, the information of $D(q(z), p(z))$ that depends on inputs of the whole batch is leaked. By this leaked sensitive information, we might get beautiful results, but it is the result of our poor understanding of gradient constructions in the DP-SGD manner.

3.3 Augmentation for Estimating Reconstruction Error

Back to the original VAE [22], the stochastic gradient variational Bayes (SGVB) estimator enables us to compute the ELBO over a single batch as:

$$\mathcal{L}_{ELBO}(x_i) = -\mathcal{L}_i = -\text{kld}_i + \frac{1}{L} \sum_{l=1}^{L} \log p(x_i|z_{i,l}) \tag{12}$$

In the original VAE, we can set $L = 1$ if the batch size is large enough [2]. However, DP-SGD assumes *micro-batches* whose size is 1. In order to accurately estimate the log-likelihood around x_i, we should set L in no small number. Thanks to gradient clipping (3), the sensitivity is still bounded by C even when utilizing a large L. Since $\frac{1}{L} \sum_{l \in [L]} \log p(x_i|z_{i,l})$ is independent from $x_{j \neq i}$, and the stochastic gradient including it is clipped by the constant C, the sensitivity is bounded by C against any L.

From the above discussion, we can utilize augmentations that reduce the reconstruction error without increasing the sensitivity. However, it consumes much more computational time and memory spaces.

4 Proposed Method

Based on the sensitivity analysis, we present how to learn differentially private variational autoencoders with suppressing the amount of noise. We first introduce a general model PRIVAE that learns variational autoencoder in a differentially private way. Second, we propose a novel learning technique *term-wise DP-SGD* that reduces the amount of noise for DP by decomposing stochastic gradients into term-wise components. Our proposed method also utilizes the augmentation that attempts to reduce the reconstruction error, as discussed in section 3.3.

4.1 PRIVAE: a general model of differentially private VAE

Our basic idea is to decompose the terms of the loss function into two groups and compose a noisy gradient that ensures the DP group by group. For each group, we separately run the gradient aggregation sequence for DP, which consists of computing stochastic gradients, clipping gradients, and adding noise as following the DP-SGD manner.

Towards reducing the amount of noise, we first introduce the notation of partitions. Let s be a partition of *batch*, where $s = \{x_1, \ldots, x_{|s|}\}$, $x_i \in batch$. Any pairs of s_j and s_k ($j \neq k$) are mutually disjoint, that is $s_j \cap s_{k \neq j} = \emptyset$.

Objective function of PRIVAE. PRIVAE minimizes the objective function described below:

$$\mathcal{L} = \mathcal{L}_{sample} + \mathcal{L}_{batch} = \mathbb{E}_x[\phi(x)] + \mathbb{E}_s[\psi(s)] = \frac{1}{B} \sum_{i \in [B]} \phi(x_i) + \frac{1}{b} \sum_{j \in [b]} \psi(s_j) \tag{13}$$

where $\phi(x) = \text{rec}(x) + \beta\text{kld}(x)$, $\psi(s) = \alpha D(q(z), p(z))$ and s is the partition denoted the above. Let b is the number of the partitions, and $\bigcup_{j \in [b]} s_j = batch$. Note, (13) with $b = 1$ is identical to (8).

---

[2][22] mentioned that L can be set to 1 as long as the minibatch size was large enough. e.g. $B = 100$.

Algorithm 1 Termwise DP-SGD

Input: $x_1, \ldots, x_N$
Parameters: learning rate η_t, noise scale σ'_ε, batch size B, #partitions b, clipping size C_1 and C_2
Initialize θ_0 randomly
for t **in** $[T]$ **do**
    randomly sample *batch* with probability B/N
    $\bar{\mathbf{g}}_{\text{sample}} \leftarrow \sum_{i \in [B]} \pi_{C_1}(\nabla_{\theta_t}\phi(x_i; \theta_t))$
    randomly sample sub-group s_j from *batch* with probability b/B
    $\bar{\mathbf{g}}_{\text{batch}} \leftarrow \sum_{j \in [b]} \pi_{C_2}(\nabla_{\theta_t}\psi(s_j; \theta_t))$
    $\tilde{\mathbf{g}} \leftarrow \frac{1}{B}(\bar{\mathbf{g}}_{\text{sample}} + \mathcal{N}(0, (\sigma'_\varepsilon C_1)^2\mathbf{I})) + \frac{1}{b}(\bar{\mathbf{g}}_{\text{batch}} + \mathcal{N}(0, (\sigma'_\varepsilon C_2)^2\mathbf{I}))$
    $\theta_{t+1} \leftarrow \theta_t - \eta_t\tilde{\mathbf{g}}$
end for
Output: θ_T

4.2 Termwise DP-SGD

We propose termwise DP-SGD that composes noisy gradient for DP in a term-wise way. The termwise DP-SGD crafts the noisy gradients for sample-wise terms $\phi(x_i)$ and batch-wise terms $\psi(s_j)$, separately. In the last phase of termwise DP-SGD, it combines these noisy gradients and updates parameters θ. The overall proposed procedure of termwise DP-SGD is in Algorithm 1.

Gradient aggregation for sample-wise term. For each sample-wise term $\phi(x_i)$, we craft its clipped gradient $\pi_{C_1}(\nabla_\theta\phi(x_i; \theta))$ with clip size C_1. We then aggregate the sum of the clipped gradients as follows:

$$\bar{\mathbf{g}}_{\text{sample}} = \sum_{i \in [B]} \pi_{C_1}(\nabla_\theta\phi(x_i; \theta)). \tag{14}$$

Gradient aggregation for batch-wise term. For the batch-wise terms $\psi(s_j)$, we first partition *batch* into sub-groups $s_1, \ldots, s_b$ where $b \leq B$. We then compute $\psi(s_j)$ for $j \in [b]$ and aggregate their clipped gradients with clip size C_2 as described below:

$$\bar{\mathbf{g}}_{\text{batch}} = \sum_{j \in [b]} \pi_{C_2}(\nabla_\theta\psi(s_j; \theta)). \tag{15}$$

Term-wise noise injections and concatenation. Finally we combine the above two gradients as

$$\tilde{\mathbf{g}} = \frac{1}{B}\left(\bar{\mathbf{g}}_{\text{sample}} + \mathcal{N}(0, (\sigma'_\varepsilon C_1)^2\mathbf{I})\right) + \frac{1}{b}\left(\bar{\mathbf{g}}_{\text{batch}} + \mathcal{N}(0, (\sigma'_\varepsilon C_2)^2\mathbf{I})\right). \tag{16}$$

where $\sigma'_\varepsilon = \kappa\sigma_\varepsilon$, $\kappa = \frac{\sigma_{\varepsilon/2}}{\sigma_\varepsilon} = 2\sqrt{\frac{\log \delta/2}{\log \delta}}$ if $C_2 > 0$, otherwise $\kappa = 1$. κ can be derived from (5). The ratio of C_2 to C_1 plays an important role that adjusts the scale of (15) against (14) like α.

4.3 Discussion

We here discuss the privacy guarantee and noise scale of our proposed method.

Theorem 2 *Term-wise DP-SGD with the noise scale σ'_ε satisfies (ε, δ)-differential privacy if DP-SGD with σ_ε satisfies (ε, δ)-differential privacy for a VAE that has no batch-wise terms.*

Proof *$\sigma'_\varepsilon = \kappa\sigma_\varepsilon$ is the noise scale that satisfies ($\varepsilon/2$, $\delta/2$)-DP. From the sequential composition of the first term and the second term in (16), the sum of the two terms satisfies (ε, δ)-DP.*

Lemma 2 *ℓ_2-sensitivity of $\mathcal{L}$ (13) is $C_1 + C_2$. That means the sensitivity is $O(1)$.*

Proof *Since all s_j and $s_{k \neq j}$ are disjoint, the change of any single $x_i \in batch$ influences only $\phi(x_i)$ and $\psi(s_j)$ where $x_i \in s_j$. Thus, ℓ_2-sensitivity of $\mathcal{L}_{sample}$ and $\mathcal{L}_{batch}$ is C_1 and C_2, respectively. Finally, ℓ_2-sensitivity of $\mathcal{L}$ is $C_1 + C_2$.*

In (13) and (16), the computation of ψ for each partition results in under-estimation against $\psi(batch)$, but it brings the reduction of the noise variance for the second term. In (16), the noise $\mathcal{N}(0, (\sigma'_\varepsilon C_2)^2\mathbf{I})$

Table 1: Sensitivity and noise scale for learning differentially private models with a batch-wise term.

	$\mathcal{L}$	ℓ_2-sensitivity of $\mathcal{L}$	noise scale
DP-SGD (micro agg.)	$\frac{1}{B}\sum_{i=1}^{B}(\phi(x_i) + \psi(batch))$	BC	$C\sigma_\varepsilon$
DP-SGD (batch agg.)	$\psi(batch)$	C	$C\sigma_\varepsilon$
Term-wise DP-SGD	$\frac{1}{B}\sum_{i=1}^{B}\phi(x_i) + \frac{1}{b}\sum_{j=1}^{b}\psi(s_j)$	$C_1 + C_2$	$(\frac{C_1}{B} + \frac{C_2}{b})\kappa\sigma_\varepsilon$

(a) Sparsity   (b) Log-likelihood   (c) MMD($q_\phi(z), p(z)$)
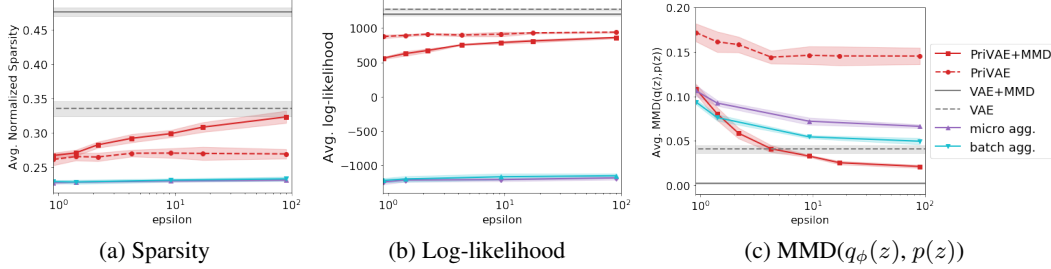
Figure 1: Our model PRIVAE works with MMD as a regularization term, and obtains sparsity.

can be divided by the number of partitions b. Therefore, we can manipulate the degree of the trade-off between the estimation accuracy of $\psi(batch)$ and the second term's noise scale by b.

Finally, we discuss the noise scale. In the existing method DP-SGD with a divergence, the overall noise scale is $\Sigma_\varepsilon = \frac{BC\sigma_\varepsilon}{B} = C\sigma_\varepsilon$. While our term-wise DP-SGD has $\Sigma'_\varepsilon = \frac{C_1\sigma'_\varepsilon}{B} + \frac{C_2\sigma'_\varepsilon}{b} \approx 2\left(\frac{C_1}{B} + \frac{C_2}{b}\right)\sigma_\varepsilon$ by using $\sigma'_\varepsilon = \kappa\sigma_\varepsilon \approx 2\sigma_\varepsilon$. In the DP-SGD with divergence, the order of the noise scale can be written as $O(1)$, while our proposed method has $O(1/b)$ since $B \geq b$.

Table 1 summarizes the sensitivity and noise scale of DP-SGD and our term-wise DP-SGD.

5 Evaluation

In this section, we demonstrate the effectiveness of our proposed method PRIVAE with two different tasks. We evaluate our method in a sparse coding task and a clustering task. Each task employs a different prior distribution as $p(z)$ and divergence as the regularization term $\psi(s)$. The experimental settings, including datasets, neural network architectures, construction of prior distributions, regularization divergences, and evaluation metrics, follow the experiments in [24]. The experimental codes are developed in Python 3.7 and PyTorch 1.5 [28] and run on machines with a Tesla V100 GPU.

5.1 Sparsity

We first consider a sparse representation that only a small fraction of available factors are employed for reconstructions. In this task, we utilize the Fashion-MNIST dataset [32]. As well as [24], we construct a sparse prior as $p(z) = \Pi_d(1 - \gamma)\mathcal{N}(z_d; 0, 1) + \gamma\mathcal{N}(z_d; 0, \sigma_0^2)$ with $\sigma_0^2 = 0.05$. This mixture distribution can be interpreted as a mixture of samples being either *off* or *on*, whose proportion is set by γ. We set γ=0.8. The regularization term we utilize here is a dimension-wise MMD with a sum of Cauchy kernels on each dimension $(k(\boldsymbol{x}, \boldsymbol{y}) = \sum_{d=1}^{D}\sum_{\ell}^{L}\frac{\sigma_\ell}{\sigma_{\ell=1}+(x_d-y_d)^2})$ with $\sigma_\ell \in \{0.2, 0.4, 1, 2, 4, 10\}$. To measure a sparsity of the latent representations, we employ the sparsity metric defined with the Hoyer extrinsic metric [18] as follows:

$$\text{Sparsity} = \frac{1}{n}\sum_{i \in [n]}\text{Hoyer}(\bar{\boldsymbol{z}}_i), \quad \text{Hoyer}(\boldsymbol{y}) = \frac{\sqrt{d} - \|\boldsymbol{y}\|_1/\|\boldsymbol{y}\|_2}{\sqrt{d} - 1} \tag{17}$$

where $\bar{\boldsymbol{z}}_i$ is a vector whose d-th dimensional value $\bar{z}_{i,d} = z_{i,d}/\sigma(z_{i,d})$. $\sigma(z_{i,d})$ is the standard deviation of d-th dimentional latent encoding taken over the dataset. The Hoyer($\boldsymbol{y}$) $\in [0, 1]$ represents 0 for fully dense vector and 1 for a fully sparse vector.

We use the same convolutional neural networks for both the encoder and decoder as in [24] with D=50 dimensional latent space. In this task, we use SGD optimizer with C_1=0.05, η=0.001, β=1,

7

(a) PRIVAE ($\varepsilon = 2.87$)   (b) PRIVAE with KL($p(z)||q(z)$) ($\varepsilon = 2.87$)
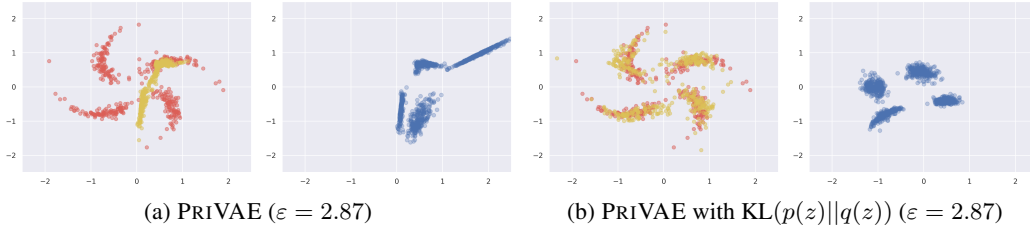
Figure 2: PRIVAE works for reconstructions with clustering representations. (a) and (b) show reconstructions (yellow) of the pinwheel data (red) and the learned representations (blue) for PRIVAE and PRIVAE with KL($p(z)||q(z)$), respectively. PRIVAE with KL($p(z)||q(z)$) demonstrates good reconstructions and representations for clustered pinweels, but PRIVAE without it shows poor results.

B=256, b=16, L=1 for all privatized models, and C_2=0.005 for PRIVAE with the MMD and C_2=0 for PRIVAE without it. For non-private VAEs, we use Adam optimizer with η=0.0005, B=256. For both VAE and PRIVAE, we set α=100 when attaching the MMD. We also compare with DP-SGD using micro-agg. and batch-agg.. For these methods, we set C=0.0002 to avoid exploding gradients. The other hyper-parameters are the same as PRIVAE with MMD. All models are trained in 10 epochs.

Figure 1 shows the substantial sparsity by the sparse prior (Figure 1a), the log-likelihood (Figure 1b), and the MMD between q(z) and p(z) (Figure 1c), those results are observed at several privacy parameter ε. We plot the average over ten observations. The shaded regions are $\pm$ 1 standard deviation around the averages. In Figure 1a, PRIVAE with the regularization (PRIVAE +MMD) demonstrates higher sparsity than the model that does not have it. Although it has a gap between the non-private regularized model (VAE+MMD), our proposed model achieved increasing the sparsity even under differential privacy constraints. In the MMD between $q(z)$ and $p(z)$, PRIVAE +MMD shows smaller values against PRIVAE without it. By employing the regularization term, PRIVAE could obtain the sparsity and reduce the MMD, but it was not easy to simultaneously increase the log-likelihood. The trade-off between them seems more significant than non-private models. To obtain more sparsity, PRIVAE needs to improve reconstruction performance.

5.2 Clustering Latent Space

Next, we consider a differentially private VAE that wishes to impose *clustering* of the latent space. For this experiment, we utilize the pinwheel dataset from [19], with n=400 observations, clustered in 4 spirals. Following the experiment in [24], we utilize a mixture of four Gaussians as the prior, KL($p(z)||q(z)$) as a regularization divergence, and fully-connected neural networks for both encoder and decoder. The prior is defined as $p(z) = \sum_{k=1}^{K} \pi^k \prod_{d=1}^{D} \mathcal{N}(\boldsymbol{z}|\mu_d^k, \sigma_d^k)$ with D=2, K=4, σ_d^k=0.03, π^k=1/K, and $\mu_d^k \in \{0, 1\}$. The divergence is defined as KL($p(z)||q(z)$) $\approx \sum_{j=1}^{|s|} (\log p(z_j) - \log \sum_{i=1}^{|s|} q(z_j|x_i))$. We set C_1=0.05, η=0.01, B=20, b=1, L=20 for all models, C_2=0.0005, β=0 for PRIVAE with KL($p(z)||q(z)$) and C_2=0, β=1 for PRIVAE without it.

We compare the clustering performance between PRIVAE with/out the regularization term KL($p(z)||q(z)$). Figure 2 shows the reconstructions of the pinwheel data and the (clustered) representations. The first two columns demonstrate the results of PRIVAE without KL($p(z)||q(z)$), and the others show those of PRIVAE with KL($p(z)||q(z)$). In the figures the red dots represent the original inputs, the yellow dots are their reconstructions, and the blue dots show the data points in the latent spaces. PRIVAE without the regularization demonstrates poor reconstructions against the raw pinwheel clustered data. While, PRIVAE with KL($p(z)||q(z)$) generated better reconstructions than the model without it even though the generated samples have small reconstruction errors. The learned representations of PRIVAE with the regularization are well clustered and fitted to the prior that is the four mixture of Gaussians. Through these results, our proposed model worked well with employing the prior and the regularization term those intended to capture the clusters of the pinwheel data.

6 Conclusion

This paper studied how to learn variational autoencoders with various divergence under differential privacy constraints. We revealed several divergences increase the sensitivity of the stochastic gradient from $O(1)$ to $O(B)$ in terms of batch size B. To reduce the sensitivity and the amount of noise, we proposed a term-wise DP-SGD that crafted randomized gradients in two different ways tailored to the compositions of the loss terms. The term-wise DP-SGD could keep the sensitivity at $O(1)$ even when attaching the divergence. In our experiments, we demonstrated that our method worked well with two pairs of the prior distribution and the divergence. We mainly studied differentially private VAEs, but these contributions also have significant importance for the other machine learning models required to satisfy differential privacy.

References

[1] Tensorflow privacy. https://github.com/tensorflow/privacy.

[2] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318. ACM, 2016.

[3] J. M. Abowd. The us census bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2867–2867, 2018.

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

[5] A. A. Alemi, I. Fischer, J. V. Dillon, and K. Murphy. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016.

[6] K. Amin, T. Dick, A. Kulesza, A. Munoz, and S. Vassilvitskii. Differentially private covariance estimation. In *Advances in Neural Information Processing Systems*, pages 14190–14199, 2019.

[7] E. Bagdasaryan, O. Poursaeed, and V. Shmatikov. Differential privacy has disparate impact on model accuracy. In *Advances in Neural Information Processing Systems*, pages 15453–15462, 2019.

[8] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.

[9] V. Bindschaedler and R. Shokri. Synthesizing plausible privacy-preserving location traces. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 546–563. IEEE, 2016.

[10] V. Bindschaedler, R. Shokri, and C. A. Gunter. Plausible deniability for privacy-preserving data synthesis. *Proceedings of the VLDB Endowment*, 10(5):481–492, 2017.

[11] K. Chaudhuri, J. Imola, and A. Machanavajjhala. Capacity bounded differential privacy. In *Advances in Neural Information Processing Systems*, pages 3469–3478, 2019.

[12] R. Chen, Q. Xiao, Y. Zhang, and J. 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. ACM, 2015.

[13] C. Dwork. Differential privacy. In *Proceedings of the 33rd international conference on Automata, Languages and Programming-Volume Part II*, pages 1–12. Springer-Verlag, 2006.

[14] C. Dwork. Differential privacy. *Encyclopedia of Cryptography and Security*, pages 338–340, 2011.

[15] C. Dwork. A firm foundation for private data analysis. *Communications of the ACM*, 54(1):86–95, 2011.

[16] C. Eastwood and C. K. Williams. A framework for the quantitative evaluation of disentangled representations. In *International Conference on Learning Representations*, 2018.

[17] B. Esmaeili, H. Wu, S. Jain, A. Bozkurt, N. Siddharth, B. Paige, D. H. Brooks, J. Dy, and J.-W. Meent. Structured disentangled representations. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2525–2534, 2019.

[18] N. Hurley and S. Rickard. Comparing measures of sparsity. *IEEE Transactions on Information Theory*, 55(10):4723–4741, 2009.

[19] M. J. Johnson, D. K. Duvenaud, A. Wiltschko, R. P. Adams, and S. R. Datta. Composing graphical models with neural networks for structured representations and fast inference. In *Advances in neural information processing systems*, pages 2946–2954, 2016.

[20] J. Jordon, J. Yoon, and M. van der Schaar. Pate-gan: generating synthetic data with differential privacy guarantees. 2018.

[21] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling. Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems*, pages 4743–4751, 2016.

[22] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[23] J. Lee and D. Kifer. Concentrated differentially private gradient descent with adaptive per-iteration privacy budget. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1656–1665, 2018.

[24] E. Mathieu, T. Rainforth, N. Siddharth, and Y. W. Teh. Disentangling disentanglement in variational autoencoders. In *International Conference on Machine Learning*, pages 4402–4412, 2019.

[25] H. B. McMahan, G. Andrew, U. Erlingsson, S. Chien, I. Mironov, N. Papernot, and P. Kairouz. A general approach to adding differential privacy to iterative training procedures. *arXiv preprint arXiv:1812.06210*, 2018.

[26] H. B. McMahan, D. Ramage, K. Talwar, and L. Zhang. Learning differentially private recurrent language models. *arXiv preprint arXiv:1710.06963*, 2017.

[27] N. Papernot, M. Abadi, U. Erlingsson, I. Goodfellow, and K. Talwar. Semi-supervised knowledge transfer for deep learning from private training data. *arXiv preprint arXiv:1610.05755*, 2016.

[28] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. 2017.

[29] A. Schein, Z. S. Wu, A. Schofield, M. Zhou, and H. Wallach. Locally private bayesian inference for count models. In *International Conference on Machine Learning*, pages 5638–5648, 2019.

[30] R. Torkzadehmahani, P. Kairouz, and B. Paten. Dp-cgan: Differentially private synthetic data and label generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.

[31] D. Wang and J. Xu. On sparse linear regression in the local differential privacy model. In *International Conference on Machine Learning*, pages 6628–6637, 2019.

[32] H. Xiao, K. Rasul, and R. Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

[33] L. Xie, K. Lin, S. Wang, F. Wang, and J. Zhou. Differentially private generative adversarial network. *arXiv preprint arXiv:1802.06739*, 2018.

[34] L. Yu, L. Liu, C. Pu, M. E. Gursoy, and S. Truex. Differentially private model publishing for deep learning. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 332–349. IEEE, 2019.

[35] J. Zhang, G. Cormode, C. M. Procopiuc, D. Srivastava, and X. Xiao. Privbayes: private data release via bayesian networks. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, pages 1423–1434, 2014.

[36] J. Zhang, X. Xiao, and X. Xie. Privtree: A differentially private algorithm for hierarchical decompositions. In *Proceedings of the 2016 International Conference on Management of Data*, pages 155–170. ACM, 2016.