
FlexAE: Flexibly Learning Latent Priors for Wasserstein Auto-Encoders

Arnab Kumar Mondal¹

Himanshu Asnani²

Parag Singla¹

Prathosh AP¹

¹IIT, Delhi, India

²TIFR, Mumbai, India

Abstract

Auto-Encoder (AE) based neural generative frameworks model the joint-distribution between the data and the latent space using an Encoder-Decoder pair, with regularization imposed in terms of a prior over the latent space. Despite their advantages, such as stability in training, efficient inference, the performance of AE based models has not reached the superior standards of the other generative models such as Generative Adversarial Networks (GANs). Motivated by this, we examine the effect of the latent prior on the generation quality of deterministic AE models in this paper. Specifically, we consider the class of Generative AE models with deterministic Encoder-Decoder pair (such as Wasserstein Auto-Encoder (WAE), Adversarial Auto-Encoder (AAE)), and show that having a fixed prior distribution, *a priori*, oblivious to the dimensionality of the ‘true’ latent space, will lead to the infeasibility of the optimization problem considered. As a remedy to the issue mentioned above, we introduce an additional state space in the form of flexibly learnable latent priors, in the optimization objective of WAE/AAE. Additionally, we employ a latent-space interpolation based smoothing scheme to address the non-smoothness that may arise from highly flexible priors. We show the efficacy of our proposed models, called FlexAE and FlexAE-SR, through several experiments on multiple datasets, and demonstrate that FlexAE-SR is the new state-of-the-art for the AE based generative models in terms of generation quality as measured by several metrics such as Fréchet Inception Distance, Precision/Recall score. Code for our paper is available at: <https://github.com/dair-iitd/FlexAE>

1 INTRODUCTION

Auto-Encoder (AE) based latent variable generative models implicitly define a joint distribution over the input data and a lower-dimensional latent space, by approximating the true latent posterior, with a variational distribution. This variational distribution is parameterized using a neural network called the Encoder. The distribution induced by the Encoder is regularized to follow a pre-defined latent prior distribution. Subsequently, a Decoder network is trained to conditionally sample from the data distribution via optimizing a data-reconstruction metric. The parameters of the Encoder and the Decoder networks are learnt by optimizing either a bound on the data likelihood [Kingma and Welling, 2014] or a divergence measure between the true and generated data distributions [Tolstikhin et al., 2018]. The framework of AE-based generative models is attractive because of its ease and stability in training, efficiency in sampling, and flexibility in architectural choices. However, despite their advantages, AE-based models have failed to reach the performance of other State-of-The-Art (SoTA) generative models [Dai and Wipf, 2019, Mondal et al., 2020].

Several aspects such as the loss function used for optimization [Higgins et al., 2017, Larsen et al., 2016], presence of conflicting terms in the optimization objective [Hoffman and Johnson, 2016, Kim and Mnih, 2018], distributional choices (e.g., Gaussianity) imposed on the Encoder and Decoder [Zhao et al., 2019, Rezende and Viola, 2018], dimensionality of the latent space used [Dai and Wipf, 2019, Mondal et al., 2020], the mismatch between the learned and imposed prior [Shengjia Zhao and Ermon, 2017, Tomczak and Welling, 2018] have been identified as possible causes for the sub-optimal performance of the AE-based models. Many remedial measures, including the modification of the objective function [Zhao et al., 2019, Higgins et al., 2017, Kim and Mnih, 2018], use of non-Gaussian Encoder/Decoder [Larsen et al., 2016, Nalisnick et al., 2016], masking of spurious latent dimensions [Mondal et al., 2020], incorporating a richer class of priors on the latent space [Tomczak and

Welling, 2018, Takahashi et al., 2019, Klushyn et al., 2019], have been proposed in the literature to address some of these issues. While these modifications have improved AE models' performance, they are still behind SoTA generative models [Dai and Wipf, 2019, Mondal et al., 2020]. In this work, we make the following contributions:

1. We theoretically establish that in a deterministic AE based generative model (such Wasserstein AE), choosing a latent prior distribution supported on the entire space, leads to an infeasible optimization objective, when the model's latent space has dimensionality that is other than that of the 'true' latent space.
2. As a remedy, we propose a new model called FlexAE, that can impose flexible learnable priors on WAEs to make the optimization problem feasible by introducing an additional state space over the latent prior.
3. We employ an adversarial regularization technique to smooth the latent space of our model with flexible priors to prevent memorization. Furthermore, We propose two novel metrics, called, Pixel Memorization Score and Inception Memorization Score, to quantitatively evaluate whether the generative AE has memorized the training samples.
4. We empirically demonstrate our claims through extensive experimentation on synthetic and real-world datasets by achieving significant improvement in generation quality over the SoTA AE-based generative models.

2 BACKGROUND

The general theme in majority of the AE based generative frameworks is to implicitly learn the joint distribution between the observed data and a latent variable, via optimizing an objective function which consists of an auto-encoding (conditional likelihood) and a latent regularization term (divergence measure). Variational Autoencoder (VAE) [Kingma and Welling, 2014] is the pioneering member of this family, in which the variational latent posterior and conditional data likelihood are respectively parameterized by probabilistic (Gaussian) Encoder and Decoder networks, while the latent prior is assumed to be an isotropic Gaussian distribution. A related class of AE-models are the Adversarial Auto-Encoders (AAEs) [Makhzani et al., 2016] and Wasserstein Auto-Encoders (WAEs) [Tolstikhin et al., 2018] where a pair of deterministic Encoder-Decoder is used with Jensen-Shannon and Wasserstein distance respectively, between the aggregated latent posterior and the latent prior.

Although VAEs/WAEs provide solid frameworks for AE-based generative models, several drawbacks are associated with them, which prevents them to compete with the other SoTA generative models. It is shown that there exists a conflict between the two terms of the objective, in the case of

VAEs [Higgins et al., 2017, Shengjia Zhao and Ermon, 2017, Rezende and Viola, 2018]. A few remedial measures such as introduction of a tunable parameter in the objective [Burgess et al., 2017], use of additional penalties such as mutual information [Zhao et al., 2019], total correlation [Kim and Mnih, 2018], and generalised optimization objective [Rezende and Viola, 2018] have been proposed. Another often discussed issue with AE-models with stochastic Encoder-Decoders is that they adopt a simple unimodal Gaussian distribution for parameterization [Rosca et al., 2018]. To address this, [Nalisnick and Smyth, 2017] implements a Bayesian non-parametric version of the variational autoencoder that has a latent representation with stochastic dimensionality and could represent richer class of distributions. Invertible flow-based generative models [Kingma et al., 2016, Rezende and Mohamed, 2015] capitalize on the idea of normalizing flow for the Encoder and Decoder networks. VAE-GAN [Larsen et al., 2016], VGH/VGH++ [Rosca et al., 2018] incorporates an adversarial learning at the Decoder so that it can represent a rich class of distributions.

Further, it is observed that there is a mismatch between the aggregated variational posterior and the latent prior, leading to sub-optimality of the divergence term in the objective and in turn poor generation [Tomczak and Welling, 2018, Dai and Wipf, 2019]. Several methods try to alleviate this problem, broadly in two ways (i) using a richer class of parametric priors on the latent space [Tomczak and Welling, 2018, Klushyn et al., 2019, Kumar et al., 2020] and (ii) using a post-hoc technique to minimize the divergence or sample from the latent space without regularizing it [Bauer and Mnih, 2019, Ghosh et al., 2020, Takahashi et al., 2019]. Among the first category of methods, VampPrior [Tomczak and Welling, 2018] assumes the prior to be a mixture of the conditional posteriors with a set of learnable pseudo-inputs. Klushyn et al. [2019] adapt the constrained optimization setting in [Rezende and Viola, 2018] and substitute the standard normal prior with a hierarchical prior and use an importance-weighted bound as the optimization objective. Huang et al. [2017], Kumar et al. [2020] learn the latent priors using normalizing flow based methods. Within the second category of methods, Bauer and Mnih [2019] learn to sample from a rich class of priors by multiplying a simplistic prior distribution with a learned acceptance function. Takahashi et al. [2019] used kernel density trick for matching the prior to the aggregated posterior. RAE-GMM [Ghosh et al., 2020], imposes an L2-norm penalty in the latent space and learns to sample from it using a Gaussian Mixture Model (GMM) on the latent space. While these methods report improvement over the SoTA metrics, not many give a theoretical justification for using richer-class of latent priors. Further, post-hoc latent samplers such as RAE-GMM do not have control over the amount of bias imposed (other than a simple objective scaling factor), that might lead to over/under fitting.

However, it has been both theoretically and empirically ob-

served that dimensionality of the latent space used has a critical impact on the performance of these models [Mondal et al., 2020, Dai and Wipf, 2019, Rubenstein et al., 2018]. Dai and Wipf [2019] study the implication of the mismatch between the dimensionality of the data and the true latent space and the role of Decoder variance, in the case of AEs with stochastic Encoders. They argue a learnable variance in the Decoder would make the objective reach negative infinity even when the aggregated posterior would not match the standard Gaussian prior not because of simplistic modelling assumption but because of mismatch between data dimensionality and the true latent dimensionality. To resolve this issue they introduce a second-stage VAE, which is used on the latent space of the first stage (which is a usual VAE), where the data and the latent dimensions match. In MaskAAE [Mondal et al., 2020], the authors noted that the generation quality degrades when there is a mismatch between the dimensionality of the true and the assumed latent space of a deterministic AE. They develop a procedure to explicitly zero-out (mask) the spurious latent dimensions via a learnable masking layer. In this backdrop, to the best of our knowledge, ours is the first study to propose a flexibly learnable prior scheme on deterministic AEs as a solution to the problem of infeasibility of the objective function.

3 PROPOSED METHOD

3.1 OPTIMALITY OF THE LATENT SPACE OF WASSERSTEIN AE

We start by assuming that the true data is generated in nature via a two-step process. First, the true latent variables are sampled from an n -dimensional space, $\tilde{\mathcal{Z}}$ according to some continuous distribution in \mathbb{R}^n . Next, a non-linear function, $f : \tilde{\mathcal{Z}} \rightarrow \mathcal{X}$ maps the true latent space, $\tilde{\mathcal{Z}}$ to the observed data space, $\mathcal{X} \subseteq \mathbb{R}^d$, with $d \gg n$, in most practical cases. In other words, observed data \mathbf{x} lies on \mathcal{X} , an n -dimensional manifold embedded in \mathbb{R}^d . We make a benign assumption on f that it can be represented using neural networks with sigmoidal (or hyperbolic tangent, ReLU, Leaky ReLU etc.) activations to arbitrary closeness. Under this model, the data could be seen as lying in an n -dimensional manifold within \mathbb{R}^d , with an underlying ground truth distribution $P_d(\mathbf{x})$. The objective of a Generative AE such as WAE [Tolstikhin et al., 2018] or AAE¹ [Makhzani et al., 2016] is to estimate (or learn to sample from) the distribution $P_d(\mathbf{x})$, given some i.i.d. samples drawn from it. The distribution learned by the model, denoted by $P_\theta(\mathbf{x})$ is given by $P_\theta(\mathbf{x}) = \int_{\mathcal{Z}} P_\theta(\mathbf{x}|z) dP_z$, where $P_\theta(\mathbf{x}|z)$ is the distribution parameterized by a deterministic Decoder neural network $D_\theta(z)$ and $P_Z(z)$ is the latent prior defined on an m -dimensional space, \mathcal{Z} . WAEs learn to sample from the distribution $P_\theta(\mathbf{x})$ by solving the following optimization

problem:

$$\inf_{\phi, \theta} \left(\mathbb{E}_{P_d(\mathbf{x})} \mathbb{E}_{Q_\phi(z|\mathbf{x})} \left[c(\mathbf{x}, D_\theta(E_\phi(\mathbf{x}))) \right] \right) \quad (1)$$

such that $Q_\phi(z) = P_Z(z)$

Here $Q_\phi(z|\mathbf{x})$ is the variational conditional posterior, which is also parameterized by a neural network called the Encoder, $E_\phi : \mathcal{X} \rightarrow \mathcal{Z}$. When E_ϕ is deterministic, $Q_\phi(z|\mathbf{x})$ is dirac-delta for every \mathbf{x} . $Q_\phi(z) = \int_{\mathbb{R}^d} Q_\phi(z|\mathbf{x}) dP_d(\mathbf{x})$ is the aggregated posterior distribution imposed by the Encoder, $c : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+$ is any measurable cost function (such as Mean Square Error (MSE), Mean Absolute Error (MAE) or adversarial loss and so on) and $\phi \in \Phi$, $\theta \in \Theta$ are the learnable parameters of Encoder and Decoder, respectively. The constrained optimization problem in Eq. 1 translates to auto-encoding the input data with a constraint (regularizer) that the aggregated distribution imposed by the Encoder matches with a predefined latent prior distribution. The constrained optimization objective of WAE (Eq. 1) can be equivalently written as an unconstrained problem by introducing a Lagrangian:

$$D_{WAE} = \inf_{\phi, \theta} \left(\underbrace{\mathbb{E}_{P_d(\mathbf{x})} \mathbb{E}_{Q_\phi(z|\mathbf{x})} \left[c(\mathbf{x}, D_\theta(E_\phi(\mathbf{x}))) \right]}_a + \lambda \cdot \underbrace{D_Z(Q_\phi(z), P_Z(z))}_b \right) \quad (2)$$

Where λ is the Lagrange multiplier, $D_Z(\cdot)$ is any divergence measure such as Kullback-Leibler, Jensen-Shannon or Wasserstein distance, between two distributions.

Note that objective in Eq. 1 becomes feasible only when $D_Z(Q_\phi(z), P_Z(z))$ becomes zero. Equipped with these, in Theorem 1, we show that when $m > n$ (most common practical case), the optimization objective (Eq. 1) does not have a feasible solution when the prior is fixed a priori to be any distribution which is supported outside of a set of countable union of all possible n -dimensional manifolds in an m -dimensional space, denoted by \mathcal{Q}_m^n . An example for such a prior is the isotropic Gaussian distribution in \mathbb{R}^m , which is the usual choice in most models.

Theorem 1 *If $m > n$, then the regularization term in the objective function of an AE-based generative model (Eq. 2), cannot be driven to zero. That is $D_Z(Q_\phi(z), P_Z(z)) > 0$, $\forall \phi$ and for any distributional divergence D_Z when the support of $P_Z(z) \notin \mathcal{Q}_m^n$.*

The above theorem (cf. supp. for proof) asserts that it is impossible to match the aggregated latent posterior to the prior when the assumed latent dimension is more than the true latent dimension and the assumed prior has full-support.

¹AAE is a special case of WAE [Bousquet et al., 2017].

Even though a stochastic Encoder can fill the ‘extra’ dimensions with external noise, we only consider the case of deterministic AE in this work, since stochasticity can lead to problems such as conflicting objectives [Burgess et al., 2017] and non-unique solutions [Dai and Wipf, 2019].

3.2 FLEXIBLY LEARNING PRIOR: FLEXAE

Based on the discussion so far, fixing a prior makes the optimization objective infeasible and no prior leads to overfitting. To alleviate these, we propose to flexibly learn the latent prior jointly with the AE-training by introducing an additional state-space in the original optimization objective (Equation 2) as follows:

$$D_{FlexAE} = \inf_{\psi, \phi, \theta} \left(\underbrace{\mathbb{E}_{P_d(\mathbf{x})} \mathbb{E}_{Q(\mathbf{z}|\mathbf{x})} [c(\mathbf{x}, D_\theta(E_\phi(\mathbf{x})))]}_a + \underbrace{\lambda \cdot D_Z(Q_\phi(\mathbf{z})||P_\psi(\mathbf{z}))}_b \right) \quad (3)$$

where $P_\psi(\mathbf{z})$ is a learnable latent prior parameterized using a neural network called the Prior-Generator (P-GEN), G_ψ , that takes an $m' \geq n$ dimensional isotropic Gaussian distribution as the input and generates sample from an m -dimensional $P_\psi(\mathbf{z})$ (refer Fig. 1). In our model, referred to as the Flexible AE or FlexAE, P-GEN is jointly trained with the AE to alternatively minimize the divergence measure and the reconstruction terms in Eq. 3. Consequently, the output of the P-GEN forms the prior that is imposed on the latent space. Please note that $D_{FlexAE} \leq D_{WAE}$ and thus the new formulation does not harm the optimization. Next, the Corollary 1.1 (proof in the supp.) below states that the divergence measure can be brought to zero with FlexAE.

Corollary 1.1 $\forall m' \geq n$, $D_Z(Q_\phi(\mathbf{z})||P_\psi(\mathbf{z}))$ (term (b) in FlexAE objective (Eq. 3) becomes zero for optimum set of parameters.

3.2.1 Smoothing the Latent Space

It is well known that, unlike in AE models with stochastic Encoders (such as VAE), the latent space of a deterministic AEs tend to be dirac-deltas since there is no source of stochasticity beyond that is inherently present with the data [Rezende and Viola, 2018]. Consequently, a parametric prior generator network (as with FlexAE) may also ‘memorize’ the latent codes and eventually learn to generate samples very similar to those present in the training data. This defeats the purpose of a generative AE. To address this issue, we employ an adversarial smoothing regularization scheme along with the flexibly learnable prior. We adapt the adversarial regularizer proposed by Berthelot et al. [2019] to learn a smooth latent space.

Specifically, while training, first a convex combination of the encoded representations, $\mathbf{z}^{(j)}, \mathbf{z}^{(k)}$ of two randomly sampled training examples, $\mathbf{x}^{(j)}, \mathbf{x}^{(k)}$ are taken and the resulting vector is called interpolated latent, \mathbf{z}_{inp} , i.e. $\mathbf{z}_{inp}^{(i)} = \gamma \mathbf{z}^{(j)} + (1 - \gamma) \mathbf{z}^{(k)}$, where $\gamma \sim \mathcal{U}[0, 1]$. Next, the decoder is trained to generate a realistic image, $\mathbf{x}_{inp} \sim P_\theta^c(\mathbf{x}_{inp})$ using \mathbf{z}_{inp} . Here, P_θ^c represents the distribution of images generated by decoding interpolated latent codes. Mathematically, we propose to minimize a divergence metric between the true data and the data generated by the decoder using the interpolated latent space. This imposes a local smoothness on the latent space since the decoder is forced to learn to generate not only from encoded latent codes but also from its random interpolations. We call the FlexAE model with smoothing regularization as FlexAE-SR. The objective in Equation 3 is modified as follows:

$$D_{FlexAE-SR} = \inf_{\psi, \phi, \theta} \left(\underbrace{\mathbb{E}_{P_d(\mathbf{x})} \mathbb{E}_{Q(\mathbf{z}|\mathbf{x})} [c(\mathbf{x}, D_\theta(\mathbf{z}))]}_a + \lambda_1 \cdot \underbrace{D_Z(Q_\phi(\mathbf{z})||P_\psi(\mathbf{z}))}_b + \lambda_2 \cdot \underbrace{D_X(P_\theta^c(\mathbf{x}_{inp})||P_d(\mathbf{x}))}_c \right) \quad (4)$$

where, $D_X(\cdot)$ is any divergence measure such as Kullback-Leibler, Jensen-Shannon or Wasserstein distance, between the two distributions, $P_\theta(\mathbf{x}_{inp})$ and $P_d(\mathbf{x})$. Similar interpolation technique is proven to be useful for learning better representation [Verma et al., 2019] and for enhancing the performance of generative model in unsupervised few-shot image generation [Wertheimer et al., 2020].

Please note, good generation may happen even when the divergence measure D_Z is not exactly zero but $Q_\phi(\mathbf{z})$ is ‘close’ enough to $P_Z(\mathbf{z})$. However, the closer $Q_\phi(\mathbf{z})$ and $P_Z(\mathbf{z})$ are, the better is the generation (see Figure 2 and 3), which is our claim (Theorem 1 is a singularity towards this).

3.2.2 Implementation

For implementation, we use MSE for c in term (a) of Eq. 4. D_Z , in principle can be chosen to be any distributional divergence such as Kullback-Leibler divergence (KLD), Jensen-Shannon divergence (JSD), Wasserstein Distance and so on. In this work, we propose to use Wasserstein distance and utilize the principle laid in [Arjovsky et al., 2017, Gulrajani et al., 2017], to optimize the divergence measure (term (b) in Equation 4). The loss functions used for different blocks of FlexAE are as follows:

1. Likelihood Loss - Realization of term a in Eq. 4:

$$L_{AE} = \frac{1}{s} \sum_{i=1}^s \|\mathbf{x}^{(i)} - D_\theta(E_\phi(\mathbf{x}^{(i)}))\|^2 \quad (5)$$

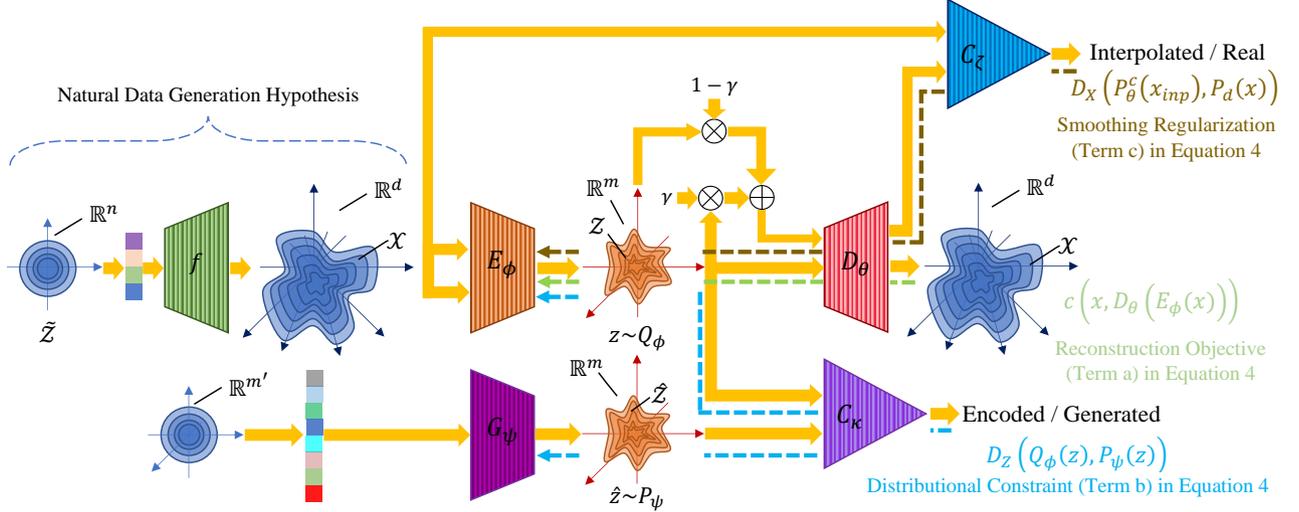


Figure 1: Proposed Model, FlexAE: Nature first samples an n -dimensional latent code from the true latent space, \tilde{z} . Next, the latent code is mapped to an n -dimensional manifold, \mathcal{X} in \mathbb{R}^d . The observed variables are encoded using a deterministic Encoder, E_ϕ . The m -dimensional encoded representations lie in an n -dimensional manifold \mathcal{Z} . The decoder network, D_θ , learns an inverse projection from the learnt latent space, \mathcal{Z} to the dataspace, \mathcal{X} . The generator network, G_ψ , parameterizes the learnable prior distribution, that maps an isotropic Gaussian distribution in $\mathbb{R}^{m'}$ to any arbitrary prior $P_\psi(z)$ in \mathbb{R}^m . The critic network, C_κ measures the distributional divergence between Q_ϕ and P_ψ . The smoothness regulator, C_ζ measures the distributional divergence between $P_\theta^c(x_{inp})$ and $P_d(x)$. The entire network is trained in an end-to-end fashion. The forward path is shown using solid yellow arrows and the flow of gradients due to different terms in the objectives are also shown as color coded dashed arrows.

2. Wasserstein Loss - We use Wasserstein distance [Arjovsky et al., 2017] for D_Z (term b Eq. 4):

$$L_{Critic} = \frac{1}{s} \sum_{i=1}^s C_\kappa(\hat{z}^{(i)}) - \frac{1}{s} \sum_{i=1}^s C_\kappa(z^{(i)}) + \frac{\beta}{s} \sum_{i=1}^s (\|\nabla_{z_{avg}^{(i)}} C_\kappa(z_{avg}^{(i)})\| - 1)^2 \quad (6)$$

$$L_{Gen} = -\frac{1}{s} \sum_{i=1}^s C_\kappa(\hat{z}^{(i)}) \quad (7)$$

$$L_{Enc} = \frac{1}{s} \sum_{i=1}^s C_\kappa(z^{(i)}) \quad (8)$$

3. Smoothing Regularization - We use Wasserstein distance [Arjovsky et al., 2017] for D_X (term c Eq. 4):

$$L_{Reg} = \frac{1}{s} \sum_{i=1}^s C_\zeta(x_{inp}^{(i)}) - \frac{1}{s} \sum_{i=1}^s C_\zeta(x^{(i)}) + \frac{\beta}{s} \sum_{i=1}^s (\|\nabla_{x_{avg}^{(i)}} C_\zeta(x_{avg}^{(i)})\| - 1)^2 \quad (9)$$

$$L_{AEReg} = -\frac{1}{s} \sum_{i=1}^s C_\zeta(x_{inp}^{(i)}) \quad (10)$$

Where, $z^{(i)} = E_\phi(x^{(i)})$, $\hat{z}^{(i)} = G_\psi(n^{(i)})$, $n^{(i)} \sim \mathcal{N}(0, I)$, $z_{inp}^{(i)} = \gamma z^{(j)} + (1 - \gamma)z^{(k)}$, $x_{inp}^{(i)} = D_\theta(z_{inp}^{(i)})$, $z_{avg}^{(i)} = \alpha z^{(i)} + (1 - \alpha)\hat{z}^{(i)}$, and $x_{avg}^{(i)} = \alpha x^{(i)} + (1 - \alpha)x_{inp}^{(i)}$. $\alpha \sim \mathcal{U}[0, 1]$, and $\beta = 10$ is a hyper-parameter as in [Gulrajani et al., 2017]. $\gamma \sim \mathcal{U}[0, 1]$. E_ϕ , D_θ , G_ψ , C_κ , and C_ζ denote the encoder, decoder, latent generator, critic, and adversarial regularization network respectively.

Finally, during inference, data is generated using FlexAE as follows:

1. Sample from a primitive (Gaussian) distribution and pass it through the P-GEN to sample a point from the latent space $P_\psi(z)$.
2. Input the latent sample through the Decoder to generate a data sample.

Algorithm for training FlexAE can be found in the supplementary material.

3.3 RELATION TO PRIOR WORKS

As discussed in section 2, a class of algorithms learn the prior to improve generation quality. Hoffman and Johnson [2016] proposed, learning the prior is one approach to optimize the ELBO in a stochastic VAE framework. Later, Huang et al. [2017] applied Real NVP [Dinh et al., 2017] to learn the prior. Tomczak and Welling [2018] proved the

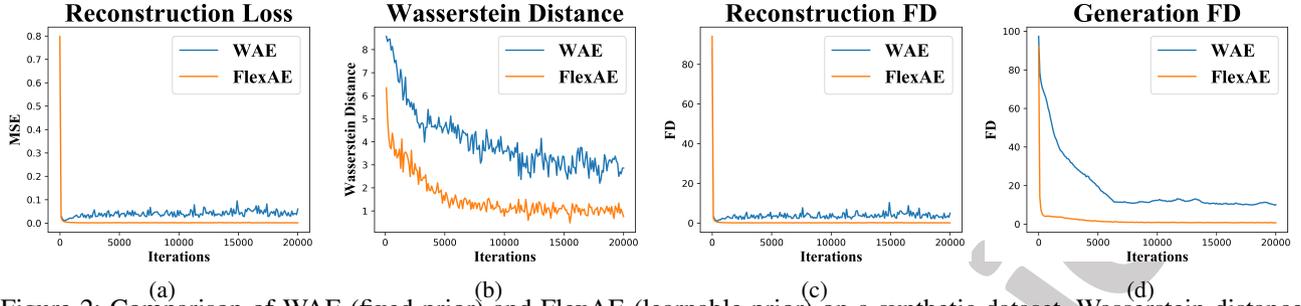


Figure 2: Comparison of WAE (fixed prior) and FlexAE (learnable prior) on a synthetic dataset. Wasserstein distance between P_z and Q_ϕ reduce faster in the case of FlexAE compared to a fixed prior WAE, leading to a better FD.

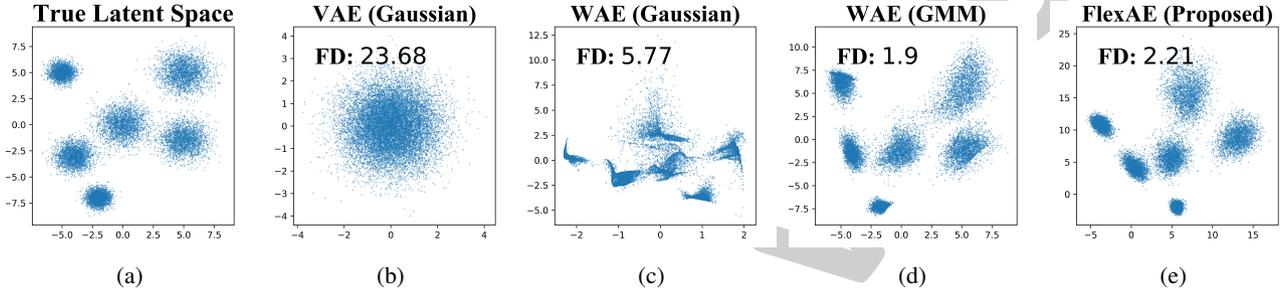


Figure 3: Visualization of (a) true latent space; (b) latent space learned by the VAE [Kingma and Welling, 2014]; (c) latent space learned by the WAE [Tolstikhin et al., 2018] with Normal prior; (d) latent space learned by the WAE [Tolstikhin et al., 2018] with GMM prior; and (e): latent space learned by the proposed FlexAE model, along with generation Fréchet Distance (FD) in each case. For multimodal data, model with multimodal prior (WAE-GMM) and FlexAE perform better.

optimal prior is the aggregated posterior, which they approximate by assembling a mixture of the posteriors with a set of learned pseudo-inputs. Bauer and Mnih [2019] proposed LARS prior, which is obtained by applying a rejection sampler with learned acceptance function to the original prior distribution. Takahashi et al. [2019] introduced the kernel density trick to estimate the KL divergence in ELBO and log-likelihood, without explicitly learning the aggregated posterior. However, Dai and Wipf [2019] argue that, if the data dimension, d mismatches with the model’s latent dimension, m , optimizing ELBO objective leads to non-unique optimum solutions. To address this concern, they train a second stage to successfully sample from the latent space of the first stage. van den Oord et al. [2017] incorporate the ideas from vector quantisation and learns a discrete latent representation paired with a learnable autoregressive prior. On the other hand, for a deterministic encoder-decoder pair, Ghosh et al. [2020], claim that an explicit l_2 regularization in the latent space is equivalent to a Gaussian prior and leads to a smooth and meaningful latent-space. For generation, they introduce an ex-post density estimation step by fitting a ten component GMM to the learned latent space.

In this work, we investigate the relation between the dimensionality of the assumed latent space and fixation of a distribution as the latent prior in generative AE models with deterministic Encoder-Decoder pair, such as WAE [Tolstikhin et al., 2018], AAE [Makhzani et al., 2016]. In this perspective, we believe that ours is the first study to the-

oretically demonstrate the infeasibility of the objective of Generative AE under a prior fixation oblivious to true latent dimension, n .

4 EXPERIMENTS AND RESULTS

4.1 SYNTHETIC EXPERIMENTS

Our objective in the synthetic experiments is to study the behaviour of loss terms, reconstruction and generation quality of FlexAE and compare them with other generative AE models.

Figure 2 demonstrates the benefit of FlexAE over WAE [Tolstikhin et al., 2018], where the performance of both the models is shown on a synthetic data: $\tilde{\mathcal{Z}} = \mathbb{R}^5$ and $f : \mathbb{R}^5 \rightarrow \mathbb{R}^{128}$ is a multi-layer perceptron with randomly initialized parameters (details in the supplementary material). It is seen that, when $m = 50$, Wasserstein distance between P_z and Q_z reduce faster and reaches much lower values in the case of FlexAE compared to a fixed prior WAE, leading to a better Fréchet Distance for generation.

Further, to demonstrate the effect of having a learnable prior, we plot the latent spaces learned using several AE models in Figure 3 on a synthetic dataset ($\tilde{\mathcal{Z}} = \mathbb{R}^2$ and $f : \mathbb{R}^2 \rightarrow \mathbb{R}^{128}$, further details in the supplementary material) where the true latent space is a mixture of Gaussian (MoG). It can be seen that the latent space learned by a FlexAE and

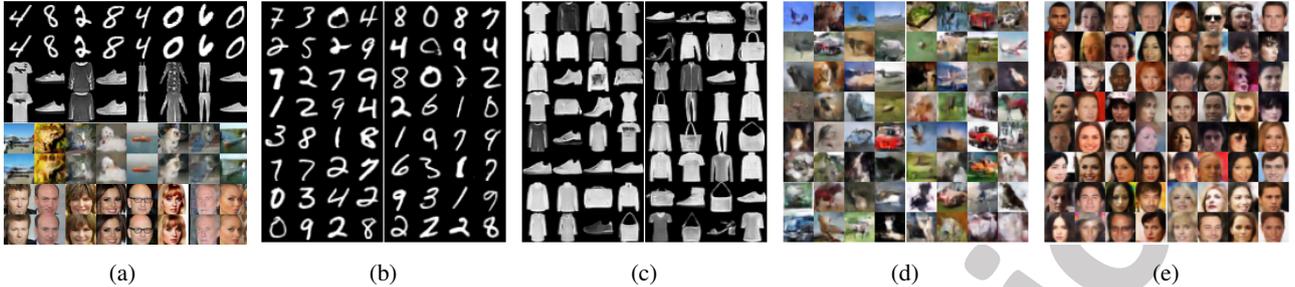


Figure 4: (a) Visualization of reconstruction quality of FlexAE-SR model on randomly selected data from the test split of MNIST (first and second rows), Fashion-MNIST (third and fourth rows), CIFAR-10 (fifth and sixth rows), and CELEBA (seventh and eighth rows). The odd rows represent the real data and the even rows represent the reconstructed data. 64 Randomly generated samples (i.e. no cherry picking) of (b) MNIST, (c) Fashion MNIST, (d) CIFAR-10, and (e) CELEBA datasets using FlexAE model shows the quality of generation of the proposed model FlexAE-SR.

Table 1: Comparison of FID scores [Heusel et al., 2017] on real datasets. Lower is better.

	MNIST		Fashion		CIFAR10		CELEBA	
	Rec.	Gen.	Rec.	Gen.	Rec.	Gen.	Rec.	Gen.
VAE [Kingma and Welling, 2014]	52.14	56.07	55.90	63.01	166.45	143.31	56.66	63.85
VAE + FLOW [Kingma et al., 2016]	28.12	33.27	34.87	49.32	90.98	123.25	39.21	45.89
VAE-VampPrior [Tomczak and Welling, 2018]	21.27	55.76	31.11	50.56	109.43	141.12	51.31	60.16
VAE-IOP [Takahashi et al., 2019]	28.01	40.16	30.20	45.39	88.17	134.72	46.51	59.35
VQ-VAE [van den Oord et al., 2017]	10.95	12.07	22.47	25.73	52.11	85.98	34.23	42.15
WAE-GAN [Tolstikhin et al., 2018]	10.48	13.60	25.93	30.21	51.36	92.30	32.49	45.58
2-S VAE [Dai and Wipf, 2019]	10.83	11.91	22.43	28.57	73.07	94.57	37.60	44.85
AE + GMM (L2) [Ghosh et al., 2020]	9.69	14.70	21.59	26.74	51.45	95.77	30.16	43.79
RAE + GMM (L2) [Ghosh et al., 2020]	9.25	10.80	19.71	25.50	50.84	90.40	34.35	44.72
MaskAAE [Mondal et al., 2020]	9.53	10.12	21.59	27.29	71.40	88.82	39.75	46.79
Pioneer [Heljakka et al., 2018]	8.17	7.30	15.91	17.25	51.07	59.12	25.35	27.94
ALAE [Pidhorskyi et al., 2020]	7.97	7.07	13.25	16.87	50.93	57.36	24.39	26.15
FlexAE (Proposed)	8.28	7.28	14.68	14.81	69.63	74.45	35.31	36.99
FlexAE-SR (Proposed)	7.63	6.75	14.17	14.41	50.36	54.61	23.94	25.78
Best GAN from Lucic et al. [2018]	-	~ 6	-	~ 20	-	~ 55	-	~ 30

a WAE with a GMM prior results in better generation as compared to the models with fixed uni-modal Gaussian priors (Note that this figure is to show that a flexible prior helps in learning but not to show impossibility).

4.2 REAL-WORLD DATASETS

We consider four real-world datasets: MNIST [Lecun, 2010], Fashion-MNIST [Xiao et al., 2017], CIFAR-10 [Krizhevsky, 2009], and CelebA [Liu et al., 2015] for our three sets of experiments. We have adopted the architecture as in 2S-VAE [Dai and Wipf, 2019] (refer to the supp.) for the AE across all the experiments for the proposed method and all baseline models to ensure a fair comparison.

4.2.1 Baseline Experiments

Methodology: The first task is to evaluate the FlexAE as a generative model. We use Fréchet Inception Distance, (FID) [Heusel et al., 2017], one of the most commonly used evaluation methods as it correlates well with human visual

perception [Lucic et al., 2018]. However, as observed in [Sajjadi et al., 2018], FID, being uni-dimensional, fails to distinguish between different cases of failure (poor sample quality and limited variation in the samples). Thus, we also report the precision and recall metrics described in [Sajjadi et al., 2018] along with FID, both of which are computed between the generated and the real test images. We compare FlexAE with a number of SoTA AE-based generative models that cover a broad class namely, VAE [Kingma and Welling, 2014], VAE+Flow [Kingma et al., 2016], VAE-VampPrior [Tomczak and Welling, 2018], VAE-IOP [Takahashi et al., 2019], VQ-VAE [van den Oord et al., 2017], WAE [Tolstikhin et al., 2018], a plain with AE post-hoc GMM, RAE+GMM [Ghosh et al., 2020], 2-stage VAE [Dai and Wipf, 2019], MaskAAE [Mondal et al., 2020], Pioneer [Heljakka et al., 2018], and ALAE [Pidhorskyi et al., 2020] with same (or similar capacity) architectures (see supplementary material). Note that a single data agnostic setting for all the hyper-parameters have been chosen for all the experiments related to FlexAE and FlexAE-SR (Refer to supplementary material for details). The details of the P-GEN and Smoothing Regularization networks can also be

Table 2: Comparison of Precision/Recall scores [Sajjadi et al., 2018] on real datasets. Higher is better.

	MNIST	Fashion	CIFAR10	CELEBA
VAE [Kingma and Welling, 2014]	0.68/0.74	0.61/0.66	0.28/0.44	0.48/0.57
2S-VAE [Dai and Wipf, 2019]	0.95/0.98	0.80/0.78	0.44/0.61	0.72/0.78
RAE + GMM (L2) [Ghosh et al., 2020]	0.96/0.97	0.78/0.77	0.49/0.59	0.67/0.75
MaskAAE [Mondal et al., 2020]	0.94/0.96	0.71/0.83	0.45/0.62	0.65/0.62
ALAE [Pidhorskyi et al., 2020]	0.98/0.99	0.97/0.96	0.73/0.85	0.84/0.81
FlexAE (Proposed)	0.98/0.99	0.98/0.98	0.65/0.80	0.71/0.76
FlexAE-SR (Proposed)	0.98/0.99	0.98/0.99	0.79/0.87	0.85/0.88

found in the supplementary.

Results: Table 1 compares the average reconstruction and generation FID scores (lower is better) of FlexAE with other AE-based generative models and the best GAN’s generation FID as reported in [Lucic et al., 2018]. It is seen that while models with parametric learnable priors (VampPrior, IOP, Flow) offer some improvement over the naive VAE, they are non optimum. It is also seen that complex prior models tend to over fit more (gap between the gen. and recon. FIDs). Further, controlling the latent space dimensionality (2SVAE, MaskAAE) have significant impact on the performance. A relatively better performance of RAE+GMM shows that while absence of prior imposition will reduce the bias, it might lead to over fitting. Pioneer [Heljakka et al., 2018] with progressively growing training procedure and adversarial encoder-generator network shows significant boost in performance². ALAE [Pidhorskyi et al., 2020] with their modified generator and discriminator architecture offers slightly better performance as compared to Pioneer. Finally, FlexAE-SR offers the best performance on three datasets as compared to other AE based generative models and their performance are comparable to that of the GANs [Lucic et al., 2018]. A similar trend is observed with the Precision/Recall scores in Table 2. FlexAE-SR achieves significantly better numbers confirming its effectiveness in generating samples that are of both high quality and variety.

Figure 4 presents reconstructed and generated samples using FlexAE-SR for qualitative evaluation of its performance. Supplementary material contains more examples.

4.2.2 Effect of Latent Space Dimensionality

Methodology: To study the effect of the latent dimensionality on the generation quality of the latent variable models, we train FlexAE-SR and WAE models with varying m .

Results: As presented in Table 3, with increasing m , the reconstruction FID decreases for both WAE and FlexAE. However, the generation FID of WAE increases with m . While generation FID of FlexAE remains almost constant.

²To ensure fairness in comparison, in this work all the methods were trained for same number of iterations.

This shows that FlexAE can achieve better optimum irrespective of the latent dimensionality.

4.2.3 Smoothness of the Latent Space

Qualitative Comparison: To ascertain the smoothness of the learned latent space and that FlexAE-SR doesn’t over fit, we conduct a few qualitative experiments: (i) traversal in the latent space between two real samples (MNIST and FMNIST) (ii) Generation by transitions in the latent space along the direction of a particular attribute (CELEBA) (iii) plot of the Nearest neighbour samples for a given generated image, from the training set, based on pixel values and inception features (CELEBA).

The outcome of these experiments are shown in Figure 5. Each row in 5a and 5b represents linear interpolation in the latent space between two randomly selected samples in the first and the last column. The central image in each row is generated using the mid-point of the two latent representations corresponding to the two real images. Each row in 5c presents manipulation of a particular face attribute (Big Nose, Heavy Makeup, Black Hair, Smiling, Male). The middle image in each row of 5c corresponds to a training sample with the attribute present. The interpolation results presented in 5a, 5b, and 5c clearly depict the smoothness of FlexAE-SR latent space as it provides smooth transition between any two random images (5a, 5b) or smooth transition based on a feature (5c). The first image of each row in 5d, and 5e shows a randomly generated sample using FlexAE-SR and the next four entries are the four nearest neighbours from the training split based on raw pixel values and inception features respectively. Visual dissimilarity between a generated image and its nearest neighbours in the training set confirms that FlexAE-SR has not merely memorized the training set. (cf. supp. for more results).

Quantitative Comparison: To quantitatively understand, if the proposed method is merely memorizing the training examples, we propose two metrics: 1) Pixel Memorization Score (PMS) and 2) Inception Memorization Score (IMS). Given two sets of images A and B, PMS computes the average L2 distance in the pixel space of all 1-nearest neighbours of images in the set A from those in the set B. IMS is a sim-

Table 3: Variation of FID w.r.t. bottleneck layer dimension, m . For MNIST and Fashion MNIST $m_b = 32$ and for CIFAR-10 and CELEBA $m_b = 64$. Unlike WAE [Tolstikhin et al., 2018], generation quality of FlexAE-SR remains unaltered even when m increases.

m	MNIST				FASHION				CIFAR10				CELEBA			
	Rec.		Gen.		Rec.		Gen.		Rec.		Gen.		Rec.		Gen.	
	WAE	FlexAE-SR	WAE	FlexAE-SR	WAE	FlexAE-SR	WAE	FlexAE-SR	WAE	FlexAE-SR	WAE	FlexAE-SR	WAE	FlexAE-SR	WAE	FlexAE-SR
m_b	7.98	6.19	15.34	5.86	20.71	10.86	40.5	11.46	51.36	50.36	92.3	54.61	32.49	23.94	45.58	25.78
$2m_b$	5.61	3.10	24.58	3.37	13.94	7.20	55.49	8.41	48.9	43.59	87.91	53.94	25.72	18.58	87.1	21.26
$4m_b$	3.99	1.64	37.16	2.79	10.04	4.93	78.31	8.16	30.47	29.71	89.82	56.09	19.89	16.21	76.12	22.18

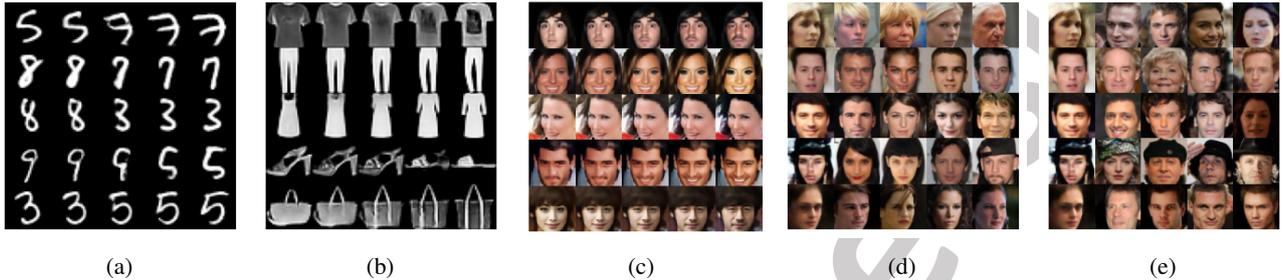


Figure 5: (a), (b) Each row in represents linear interpolation in the latent space between two randomly selected test samples in the first and the last entry. (c) Each row presents manipulation of a particular face attribute (Big Nose, Heavy Makeup, Black Hair, Smiling, Male). The central image of each row is a true image from the train split with the attribute. (d) The first image in each row shows randomly generated samples using FlexAE-SR and the next four entries are the four nearest neighbours from training data. (e) First entry in each row shows the same randomly generated samples as in (d) and the four nearest neighbours based on features extracted by a pretrained Inception net.

ilar distance metric in the inception feature space. Table 4 shows the results. We can see that, PMS between two non-overlapping sets of $10k$ training(Tr) / test(Te) images are comparable to PMS between $10k$ generated samples and $10k$ training samples. It implies that the generated samples are not replicas of the training samples seen by the model. Otherwise, PMS(Gen, Tr) would be less than PMS(Tr, Tr). Similar argument holds true for the proposed IMS metric. This shows the model has not memorized the training samples.

Table 4: Pixel/Inception Memorization Score

	Tr-Tr	Gen-Tr	Tr-Te	Gen-Te
MNIST	4.72/7.11	4.66/7.10	4.77/7.12	4.66/7.12
Fashion	4.01/9.22	3.99/9.13	4.01/9.26	3.90/9.16
CIFAR-10	9.75/12.54	9.64/14.40	9.74/12.51	9.68/14.40
CELEBA	18.42/8.73	18.01/9.37	18.50/8.75	17.86/9.34

5 CONCLUSION

In this paper, we systematically studied the effect of the latent prior on the deterministic AE-based generative models. We demonstrated that fixing any kind of prior in a data-agnostic way is detrimental to the performance. We proposed a model called the FlexAE, where we have introduced an additional state space to address the problem of infeasibility that arises due to latent dimensionality mismatch and prior fixation. We also employ a smoothing regularization technique to learn a locally convex smooth latent space for deterministic generative autoencoders. We have empirically demonstrated the efficacy of the proposed models on several real world datasets.

Author Contributions

Arnab Kumar Mondal led the work and performed all the experiments. The other authors contributed in refining the core ideas and helped in developing the theory. All the authors contributed in writing the paper.

Acknowledgements

We thank IIT Delhi HPC facility³ for computational resources. Parag Singla is supported by the DARPA Explainable Artificial Intelligence (XAI) Program with number N66001-17-2-4032, the Visvesvaraya Young Faculty Fellowships by Govt. of India and IBM SUR awards. Himanshu Asnani acknowledges the support of Department of Atomic Energy, Government of India, under the project no. 12-R&D-TFR-5.01-0500 and a gift from Adobe Research. Any opinions, findings, conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views or official policies, either expressed or implied, of the funding agencies.

References

Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *Proc. of ICML*, 2017.

³<http://supercomputing.iitd.ac.in>

- Matthias Bauer and Andriy Mnih. Resampled priors for variational autoencoders. In *Proc. of AISTATS*, 2019.
- David Berthelot, Colin Raffel, Aurko Roy, and Ian Goodfellow. Understanding and improving interpolation in autoencoders via an adversarial regularizer. In *Proc. of ICLR*, 2019.
- Olivier Bousquet, Sylvain Gelly, Ilya Tolstikhin, Carl-Johann Simon-Gabriel, and Bernhard Schoelkopf. From optimal transport to generative modeling: the vegan cookbook. *arXiv preprint arXiv:1705.07642*, 2017.
- Christopher P Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. Understanding disentangling in β -VAE. In *NeurIPS Workshop*, 2017.
- Bin Dai and David Wipf. Diagnosing and enhancing vae models. In *Proc. of ICLR*, 2019.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *Proc. of ICLR*, 2017.
- Partha Ghosh, Mehdi S. M. Sajjadi, Antonio Vergari, Michael Black, and Bernhard Scholkopf. From variational to deterministic autoencoders. In *Proc. of ICLR*, 2020.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. In *Proc. of NeurIPS*, 2017.
- Ari Heljakka, Arno Solin, and Juho Kannala. Pioneer networks: Progressively growing generative autoencoder. In *Proc. of ACCV*, 2018.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Proc. of NeurIPS*, 2017.
- Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. β -VAE: Learning basic visual concepts with a constrained variational framework. In *Proc. of ICLR*, 2017.
- Matthew D Hoffman and Matthew J Johnson. Elbo surgery: yet another way to carve up the variational evidence lower bound. In *NeurIPS Workshop*, 2016.
- Chin-Wei Huang, Ahmed Touati, Laurent Dinh, Michal Drozdal, Mohammad Havaei, Laurent Charlin, and Aaron Courville. Learnable explicit density for continuous latent space and variational inference. *arXiv preprint arXiv:1710.02248*, 2017.
- Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In *Proc. of ICML*, 2018.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *Proc. of ICLR*, 2014.
- Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. In *Proc. of NeurIPS*, 2016.
- Alexej Klushyn, Nutan Chen, Richard Kurle, Botond Cseke, and Patrick van der Smagt. Learning hierarchical priors in VAEs. In *Proc. of NeurIPS*, 2019.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- Abhishek Kumar, Ben Poole, and Kevin Murphy. Regularized autoencoders via relaxed injective probability flow. In *Proc. of AISTATS*, 2020.
- Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric. In *Proc. of ICML*, 2016.
- Y. Lecun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 2010.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proc. of ICCV*, 2015.
- Mario Lucic, Karol Kurach, Marcin Michalski, Olivier Bousquet, and Sylvain Gelly. Are gans created equal? a large-scale study. In *Proc. of NeurIPS*, 2018.
- Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, and Ian Goodfellow. Adversarial autoencoders. In *Proc. of ICLR*, 2016.
- Arnab Kumar Mondal, Sankalan Pal Chowdhury, Aravind Jayendran, Parag Singla, Himanshu Asnani, and AP Prathosh. MaskAAE: Latent space optimization for adversarial auto-encoders. In *Proc. of UAI*, 2020.
- Eric Nalisnick and Padhraic Smyth. Stick-breaking variational autoencoders. In *Proc. of ICLR*, 2017.
- Eric Nalisnick, Lars Hertel, and Padhraic Smyth. Approximate inference for deep latent gaussian mixtures. In *NeurIPS Workshop*, 2016.
- Stanislav Pidhorskyi, Donald A Adjeroh, and Gianfranco Doretto. Adversarial latent autoencoders. In *Proc. of CVPR*, 2020.
- Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *Proc. of ICML*, 2015.
- Danilo Jimenez Rezende and Fabio Viola. Taming vaes. *arXiv preprint arXiv:1810.00597*, 2018.

Mihaela Rosca, Balaji Lakshminarayanan, and Shakir Mohamed. Distribution matching in variational inference. *arXiv preprint arXiv:1802.06847*, 2018.

Paul K. Rubenstein, Bernhard Schoelkopf, and Ilya Tolstikhin. Wasserstein auto-encoders: Latent dimensionality and random encoders. In *Proc. of ICLR Workshop*, 2018.

Mehdi S. M. Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. In *Proc. of NeurIPS*, 2018.

Jiaming Song Shengjia Zhao and Stefano Ermon. Towards deeper understanding of variational autoencoding models. *arXiv preprint arXiv:1702.08658*, 2017.

Hiroshi Takahashi, Tomoharu Iwata, Yuki Yamanaka, Masanori Yamada, and Satoshi Yagi. Variational autoencoder with implicit optimal priors. In *Proc. of AAAI*, 2019.

Ilya Tolstikhin, Olivier Bousquet, Sylvain Gelly, and Bernhard Scholkopf. Wasserstein auto-encoders. In *Proc. of ICLR*, 2018.

Jakub M Tomczak and Max Welling. VAE with a vampprior. In *Proc. of AISTATS*, 2018.

Aaron van den Oord, Oriol Vinyals, et al. Neural discrete representation learning. In *Proc. of NeurIPS*, pages 6306–6315, 2017.

Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In *Proc. of ICML*, 2019.

Davis Wertheimer, Omid Poursaeed, and Bharath Hariharan. Augmentation-interpolative autoencoders for unsupervised few-shot image generation. *arXiv preprint arXiv:2011.13026*, 2020.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.

Shengjia Zhao, Jiaming Song, and Stefano Ermon. InfoVAE: Balancing learning and inference in variational autoencoders. In *Proc. of AAAI*, 2019.