

Learning Joint Latent Space EBM Prior Model for Multi-layer Generator

Jiali Cui¹, Ying Nian Wu², Tian Han¹

¹Department of Computer Science, Stevens Institute of Technology

²Department of Statistics, University of California, Los Angeles

{jcui7,than6}@stevens.edu, ywu@stat.ucla.edu

Abstract

This paper studies the fundamental problem of learning multi-layer generator models. The multi-layer generator model builds multiple layers of latent variables as a prior model on top of the generator, which benefits learning complex data distribution and hierarchical representations. However, such a prior model usually focuses on modeling inter-layer relations between latent variables by assuming non-informative (conditional) Gaussian distributions, which can be limited in model expressivity. To tackle this issue and learn more expressive prior models, we propose an energy-based model (EBM) on the joint latent space over all layers of latent variables with the multi-layer generator as its backbone. Such joint latent space EBM prior model captures the intra-layer contextual relations at each layer through layer-wise energy terms, and latent variables across different layers are jointly corrected. We develop a joint training scheme via maximum likelihood estimation (MLE), which involves Markov Chain Monte Carlo (MCMC) sampling for both prior and posterior distributions of the latent variables from different layers. To ensure efficient inference and learning, we further propose a variational training scheme where an inference model is used to amortize the costly posterior MCMC sampling. Our experiments demonstrate that the learned model can be expressive in generating high-quality images and capturing hierarchical features for better outlier detection.

1. Introduction

Deep generative models (a.k.a, *generator models*) have made promising progress in learning complex data distributions and achieved great successes in image and video synthesis [21, 34, 37, 39] as well as representation learning [5, 48]. Such models usually consist of low-dimensional latent variables together with a top-down generation model that maps such latent factors to the observed data. The latent factors can serve as an abstract data representation, but it is often modelled via a single latent vector with non-

informative prior distribution which leads to limited model expressivity and fails to capture different levels of abstractions. Learning an informative prior model for hierarchical representations is needed, yet research in this direction is still under-developed.

A principled way to learn such a prior model is by learning the generator models with multiple layers of latent variables. However, the learning of multi-layer generator model can be challenging as the inter-layer structural relation (i.e., latent variables across different layers) and the intra-layer contextual relation (i.e., latent units within the same layer) have to be effectively modelled and efficiently learned. Various methods have been proposed [5, 28, 32, 35, 40], but they only focused on inter-layer modeling by assuming the conditional Gaussian distribution across different layers while ignoring the intra-layer contextual modeling as the latent units are *conditional independent* within each layer.

The energy-based models (EBMs), on the other hand, are shown to be expressive and proved to be powerful in capturing contextual and non-structural data regularities. Notably, [33] considers the EBM in the latent space for the non-hierarchical generator model, where the energy function is considered as a correction of the non-informative Gaussian prior. The low dimensionality of the latent space makes EBM effective in capturing regularities in the data. However, a single latent vector in [33] is infeasible for capturing the patterns at multiple layers of abstractions, which limits its model capacity.

In this paper, we propose to combine the strengths of the latent space EBM and the generator with multiple layers of latent variables for better hierarchical representations and a more expressive prior model. Specifically, we introduce layer-wise energy terms to exponentially tilt the non-informative Gaussian conditional at each layer, and latent variables across different layers are modelled jointly through EBM with the multi-layer generator model as its backbone. Such a joint EBM prior model seamlessly integrates the intra-layer contextual modeling via layer-wise energy terms and inter-layer structural modeling with multi-layer latent variables.

The joint EBM prior model can be learned by maximum likelihood estimation (MLE). Each learning iteration involves Markov chain Monte Carlo (MCMC) sampling of latent variables in each layer from both the prior and posterior distributions. The prior sampling can be efficiently done due to the low dimensionality of the latent variables and, more importantly, the lightweight networks for energy functions, while the posterior sampling can be less efficient. Therefore, we further develop the variational training scheme where an additional inference model is used for posterior approximation and is jointly trained with the joint EBM prior model.

Contributions: 1) We propose a joint latent space EBM prior model for the generator model with multiple layers of latent variables; 2) We develop the maximum likelihood learning algorithm that learns the joint EBM prior model based on MCMC prior and posterior sampling across different layers. We further propose the variational joint training scheme for efficient learning and inference; 3) We provide strong empirical results through extensive experiments.

2. Background

In this section, we present the background of multi-layer latent variable model and latent space EBM prior model, which shall serve as the foundation of the proposed model.

2.1. Multi-layer latent variable model

Let \mathbf{x} be the high-dimensional observed example, and \mathbf{z} be the low-dimensional latent variables. The latent variable generative model, or *generator model*, factorizes a joint distribution of (\mathbf{x}, \mathbf{z}) as

$$p_{\beta}(\mathbf{x}, \mathbf{z}) = p_{\beta_0}(\mathbf{x}|\mathbf{z})p_{\beta_{>0}}(\mathbf{z}) \quad (1)$$

where $p_{\beta_0}(\mathbf{x}|\mathbf{z})$ is the generation model with parameter β_0 that maps from latent space to data space, and $p_{\beta_{>0}}(\mathbf{z})$ is the prior distribution over latent variables with parameter $\beta_{>0}$. $\beta = \{\beta_0, \beta_{>0}\}$.

Gaussian prior model: For non-hierarchical models [12, 23], $p_{\beta_{>0}}(\mathbf{z})$ is defined on single layer of latent variables and is typically assumed to be uniform or unit Gaussian. For hierarchical models with multiple layers of latent variables [32, 35], $p_{\beta_{>0}}(\mathbf{z})$ can be further decomposed into conditional distributions between consecutive layers of latent variables as

$$p_{\beta_{>0}}(\mathbf{z}) = \prod_{i=1}^{L-1} p_{\beta_i}(\mathbf{z}_i|\mathbf{z}_{i+1})p(\mathbf{z}_L) \quad (2)$$

where $p_{\beta_i}(\mathbf{z}_i|\mathbf{z}_{i+1}) \sim \mathcal{N}(\mu_{\beta_i}(\mathbf{z}_{i+1}), \sigma_{\beta_i}^2(\mathbf{z}_{i+1}))$ and is parameterized by a network with parameter β_i , and $p(\mathbf{z}_L)$ is chosen to be a simple distribution, such as uniform or unit Gaussian.

Maximum likelihood learning: Learning such latent variable generative models can be done using maximum likelihood estimation (MLE). The marginal distribution is $p_{\beta}(\mathbf{x}) = \int p_{\beta}(\mathbf{x}, \mathbf{z})d\mathbf{z}$ with the gradient:

$$\nabla_{\beta} \log p_{\beta}(\mathbf{x}) = \mathbb{E}_{p_{\beta}(\mathbf{z}|\mathbf{x})}[\nabla_{\beta} \log p_{\beta}(\mathbf{x}, \mathbf{z})] \quad (3)$$

where the expectation can be approximated via Monte Carlo sampling from the posterior distribution $p_{\beta}(\mathbf{z}|\mathbf{x})$. The MLE can then be accomplished through gradient ascent using such gradients. The posterior sampling usually requires the Markov Chain Monte Carlo (MCMC) such as Langevin dynamics [13, 32]

Variational learning: To alleviate the computational burden of MCMC, variational approach [4] introduces an additional inference model $q_{\omega}(\mathbf{z}|\mathbf{x})$ with a separate set of parameters $\omega = (\omega_1, \dots, \omega_L)$ for posterior approximation,

$$q_{\omega}(\mathbf{z}|\mathbf{x}) = q_{\omega_1}(\mathbf{z}_1|\mathbf{x}) \prod_{i=1}^{L-1} q_{\omega_{i+1}}(\mathbf{z}_{i+1}|\mathbf{z}_i) \quad (4)$$

where $q_{\omega_1}(\mathbf{z}_1|\mathbf{x})$ and $q_{\omega_{i+1}}(\mathbf{z}_{i+1}|\mathbf{z}_i)$ are usually assumed as conditional Gaussian distributions, forming a “bottom-up” inference structure. The generator and inference model can be jointly learned via maximizing the evidence lower bound (ELBO), i.e., $\max_{\beta, \omega} \text{ELBO}(\beta, \omega)$, where ELBO is defined as $\text{ELBO}(\beta, \omega) = \mathbb{E}_{q_{\omega}(\mathbf{z}|\mathbf{x})}[\log p_{\beta_0}(\mathbf{x}|\mathbf{z})] - D_{\text{KL}}(q_{\omega}(\mathbf{z}|\mathbf{x})||p_{\beta_{>0}}(\mathbf{z}))$.

2.2. Latent space energy-based model

The energy-based model (EBM) offers a flexible approach for learning the data distribution and is shown to be expressive in capturing data regularities [7, 8, 10, 31, 43, 45]. Most existing works focus on learning the EBM on data space, which is high-dimensional and can be challenging. To tackle this challenge, [2, 33] propose to learn latent space EBM as an informative prior model. With low-dimensional latent space, learning the EBM can be more efficient and effective, which in turn benefits the expressivity of the whole model. Specifically, [33] considers the latent space energy-based prior model on a *single layer* of latent variables,

$$p_{\alpha}(\mathbf{z}) = \frac{1}{Z(\alpha)} \exp [f_{\alpha}(\mathbf{z})]p_0(\mathbf{z}) \quad (5)$$

where $-f_{\alpha}(\mathbf{z})$ is the energy function, $Z(\alpha)$ is the normalizing constant, i.e., $Z(\alpha) = \int \exp [f_{\alpha}(\mathbf{z})]p_0(\mathbf{z})d\mathbf{z}$, and $p_0(\mathbf{z})$ is the reference distribution assumed to be unit Gaussian. Compared to data space EBMs in which the energy function needs to support the entire high-dimensional space, such exponential tilting latent space EBMs can be more efficient in capturing data regularities.

3. Model and Learning

3.1. Joint latent space EBM prior model

For generator models with multi-layer latent variables (or *multi-layer generator model*), consecutive layers are modelled by conditional Gaussian distributions (see Eqn.2), which essentially assumes the *conditional independence* for latent units within the i -th layer given the $(i + 1)$ -th layer of latent variables. Such a conditional independence assumption limits the model capacity as the contextual relation between latent units within each layer is largely ignored (see Fig.1), and needs to be improved for informative conditional modeling and better model expressivity. In this paper, we propose the joint EBM prior for multi-layer generator models,

$$\begin{aligned} p_{\alpha, \beta > 0}(\mathbf{z}) &= \frac{1}{Z_{\alpha, \beta > 0}} \exp[f_{\alpha}(\mathbf{z})] p_{\beta > 0}(\mathbf{z}) \quad (6) \\ &= \frac{1}{Z_{\alpha, \beta > 0}} \exp\left[\sum_{i=1}^L f_{\alpha_i}(\mathbf{z}_i)\right] \prod_{i=1}^{L-1} p_{\beta_i}(\mathbf{z}_i | \mathbf{z}_{i+1}) p(\mathbf{z}_L) \end{aligned}$$

where we denote $\alpha = (\alpha_1, \dots, \alpha_L)$ for EBM parameters and $\mathbf{z} = (\mathbf{z}_1, \dots, \mathbf{z}_L)$ for latent variables in different layers with layer L being the top layer. $Z_{\alpha, \beta > 0} = \int \exp[f_{\alpha}(\mathbf{z})] p_{\beta > 0}(\mathbf{z}) d\mathbf{z}$ is the normalizing constant regarding latent variables for all layers. Thus, the latent variables across different layers are *jointly* corrected via EBM prior as in Eqn.6, where $f_{\alpha}(\mathbf{z})$ is the energy function for latent variables from all layers.

In this paper, we consider a simple factorized layer-wise parameterization, i.e., $f_{\alpha}(\mathbf{z}) = \sum_{i=1}^L f_{\alpha_i}(\mathbf{z}_i)$, but other parameterizations are also feasible, which we will explore in future work. With such energy parameterization, it's worth noting that the *un-normalized* prior model can be viewed as layer-wise exponential tilting,

$$\begin{aligned} &\underbrace{\exp\left[\sum_{i=1}^L f_{\alpha_i}(\mathbf{z}_i)\right]}_{\text{Energy correction}} \underbrace{\prod_{i=1}^{L-1} p_{\beta_i}(\mathbf{z}_i | \mathbf{z}_{i+1}) p(\mathbf{z}_L)}_{\text{Gaussian prior}} \quad (7) \\ &= \underbrace{\exp[f_{\alpha_L}(\mathbf{z}_L)] p(\mathbf{z}_L)}_{\text{Correction on top layer}} \prod_{i=1}^{L-1} \underbrace{\exp[f_{\alpha_i}(\mathbf{z}_i)] p_{\beta_i}(\mathbf{z}_i | \mathbf{z}_{i+1})}_{\text{Correction on intermediate layer}} \end{aligned}$$

See Fig.1 for an illustration and comparison with multi-layer generator model with Gaussian prior.

Joint vs. conditional EBM prior: Besides the proposed joint modeling, it is also tempting to consider EBM prior for layer-wise Gaussian conditional, i.e., $\tilde{p}_{\alpha_i, \beta_i}(\mathbf{z}_i | \mathbf{z}_{i+1}) = \frac{1}{Z(\mathbf{z}_{i+1})} \exp[f_{\alpha_i}(\mathbf{z}_i)] p_{\beta_i}(\mathbf{z}_i | \mathbf{z}_{i+1})$, and form the overall prior $p_{\alpha, \beta > 0}(\mathbf{z}) = \prod_{i=1}^L \tilde{p}_{\alpha_i, \beta_i}(\mathbf{z}_i | \mathbf{z}_{i+1}) p(\mathbf{z}_L)$. Such a scheme is closely related to autoregressive energy machine

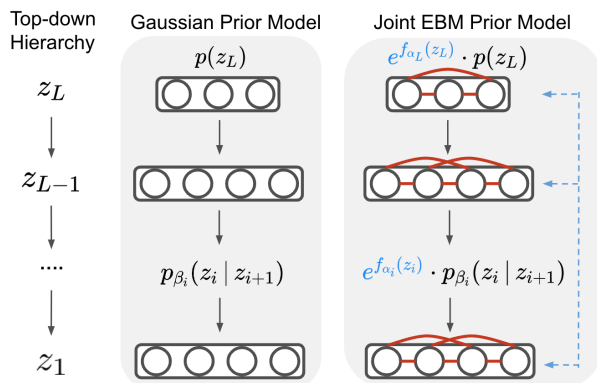


Figure 1. **Left panel:** Gaussian prior model. **Right panel:** Joint EBM prior model. **Black solid lines with arrow:** inter-layer relations modelling. **Red solid lines:** intra-layer contextual relations modelling. **Blue dashed lines:** joint modelling upon all layers.

[9] and is adopted in NCP-VAE [2]. However, the normalizing constant $Z(\mathbf{z}_{i+1})$ in $\tilde{p}_{\alpha_i, \beta_i}(\mathbf{z}_i | \mathbf{z}_{i+1})$ involves the latent variable \mathbf{z}_{i+1} from the upper layer which can be *intractable* and needs an additional inner-loop for sampling or optimization. The proposed joint EBM prior couples the latent variables across different layers via energy function and can be learned effectively and efficiently.

3.2. Maximum Likelihood Estimation

Our joint EBM prior model can be trained using MLE. Let $\theta = (\alpha, \beta)$ denotes the model parameters and θ can be learned by maximizing the log-likelihood on n training observations

$$L(\theta) = \sum_{i=1}^n \log p_{\theta}(\mathbf{x}_i) = \sum_{i=1}^n \log \int p_{\beta_0}(\mathbf{x} | \mathbf{z}) p_{\alpha, \beta > 0}(\mathbf{z}) d\mathbf{z}$$

When n becomes sufficiently large, maximizing the above log-likelihood is equivalent to minimizing the Kullback-Leibler (KL) divergence between model distribution and empirical data distribution, i.e., $\min_{\theta} D_{\text{KL}}(p_{\text{data}}(\mathbf{x}) || p_{\theta}(\mathbf{x}))$.

To update the parameter θ , we can compute the the gradient of log-likelihood $\nabla_{\theta} \log p_{\theta}(\mathbf{x})$ as

$$\begin{aligned} \nabla_{\theta} \log p_{\theta}(\mathbf{x}) &= \mathbb{E}_{p_{\theta}(\mathbf{z} | \mathbf{x})} [\nabla_{\theta} \log p_{\theta}(\mathbf{x}, \mathbf{z})] \\ &= \mathbb{E}_{p_{\theta}(\mathbf{z} | \mathbf{x})} [\nabla_{\theta} \log p_{\beta_0}(\mathbf{x} | \mathbf{z})] \quad (8) \\ &\quad + \mathbb{E}_{p_{\theta}(\mathbf{z} | \mathbf{x})} [\nabla_{\theta} \log p_{\alpha, \beta > 0}(\mathbf{z})] \end{aligned}$$

With such a gradient, we can learn θ using gradient ascent. **Learning generation model β_0 :** $p_{\beta_0}(\mathbf{x} | \mathbf{z})$ is assumed to be Gaussian distribution, i.e., $p_{\beta_0}(\mathbf{x} | \mathbf{z}) \sim \mathcal{N}(g_{\beta_0}(\mathbf{z}), \sigma^2 I)$, with generation network g_{β_0} with parameter β_0 and pre-specified σ^2 for simplicity. The learning gradient

$\nabla_{\beta_0} \log p_{\theta}(\mathbf{x})$ can then be expressed as

$$\begin{aligned} \nabla_{\beta_0} \log p_{\theta}(\mathbf{x}) &= \mathbb{E}_{p_{\theta}(\mathbf{z}|\mathbf{x})} [\nabla_{\beta_0} \log p_{\beta_0}(\mathbf{x}|\mathbf{z})] \\ &= \mathbb{E}_{p_{\theta}(\mathbf{z}|\mathbf{x})} \left[-\nabla_{\beta_0} \frac{\|\mathbf{x} - g_{\beta_0}(\mathbf{z})\|^2}{2\sigma^2} \right] \end{aligned} \quad (9)$$

Learning prior model $\alpha, \beta_{>0}$: Learning α_i can be done by computing the gradient $\nabla_{\alpha_i} \log p_{\theta}(\mathbf{x})$ as

$$\begin{aligned} \nabla_{\alpha_i} \log p_{\theta}(\mathbf{x}) &= \mathbb{E}_{p_{\theta}(\mathbf{z}|\mathbf{x})} [\nabla_{\alpha_i} f_{\alpha_i}(\mathbf{z}_i)] \\ &\quad - \mathbb{E}_{p_{\alpha, \beta_{>0}}(\mathbf{z})} [\nabla_{\alpha_i} f_{\alpha_i}(\mathbf{z}_i)] \end{aligned} \quad (10)$$

For updating $\beta_{>0}$, the gradient $\nabla_{\beta_i} \log p_{\theta}(\mathbf{x})$ is

$$\begin{aligned} \nabla_{\beta_i} \log p_{\theta}(\mathbf{x}) &= \mathbb{E}_{p_{\theta}(\mathbf{z}|\mathbf{x})} [\nabla_{\beta_i} \log p_{\beta_i}(\mathbf{z}_i|\mathbf{z}_{i+1})] \\ &\quad - \mathbb{E}_{p_{\alpha, \beta_{>0}}(\mathbf{z})} [\nabla_{\beta_i} \log p_{\beta_i}(\mathbf{z}_i|\mathbf{z}_{i+1})] \end{aligned} \quad (11)$$

Sampling: Both Eqn.10 and Eqn.11 require sampling from the posterior and prior distribution, which can be done via Langevin dynamic (LD) [26]. Given a target distribution $p(\mathbf{z})$, Langevin dynamic samples $\mathbf{z} \sim p(\mathbf{z})$ by computing the gradient $\nabla_{\mathbf{z}} \log p(\mathbf{z})$ and iteratively update \mathbf{z} as

$$\mathbf{z}_t = \mathbf{z}_{t-1} + s \nabla_{\mathbf{z}} \log p(\mathbf{z}_{t-1}) + \sqrt{2s} \epsilon_{t-1} \quad (12)$$

where t indexes the time step, s is the step size, and ϵ is the Gaussian noise for each time step.

Prior sampling: By replacing target $p(\mathbf{z})$ with $p_{\alpha, \beta_{>0}}(\mathbf{z})$, the prior sampling computes $\nabla_{\mathbf{z}} \log p_{\alpha, \beta_{>0}}(\mathbf{z})$ as

$$\nabla_{\mathbf{z}} \left[\sum_{i=1}^L f_{\alpha_i}(\mathbf{z}_i) + \sum_{i=1}^{L-1} \log p_{\beta_i}(\mathbf{z}_i|\mathbf{z}_{i+1}) + \log p(\mathbf{z}_L) \right] \quad (13)$$

Posterior sampling: By replacing $p(\mathbf{z})$ with $p_{\theta}(\mathbf{z}|\mathbf{x})$, where $p_{\theta}(\mathbf{z}|\mathbf{x}) \propto p_{\beta_0}(\mathbf{x}|\mathbf{z})p_{\alpha, \beta_{>0}}(\mathbf{z})$, the posterior sampling computes $\nabla_{\mathbf{z}} \log p_{\theta}(\mathbf{z}|\mathbf{x})$ as

$$\nabla_{\mathbf{z}} \log p_{\theta}(\mathbf{z}|\mathbf{x}) = \nabla_{\mathbf{z}} [\log p_{\beta_0}(\mathbf{x}|\mathbf{z}) + \log p_{\alpha, \beta_{>0}}(\mathbf{z})] \quad (14)$$

Notice that posterior sampling can be computationally inefficient as $\nabla_{\mathbf{z}} [\log p_{\beta_0}(\mathbf{x}|\mathbf{z})]$ requires back-propagation through the deep generation model.

3.3. Variational Learning

For efficient posterior sampling, an inference model $q_{\omega}(\mathbf{z}|\mathbf{x})$ with a separate set of parameters ω can be used for the posterior approximation. In this paper, we use the bottom-up inference model as Eqn.4 for amortizing the costly posterior MCMC sampling. Particularly, instead of KL minimization between marginal distributions as in MLE (see Sec.3.2), we consider the KL optimization between two joint densities, one for generator model density, i.e., $p_{\theta}(\mathbf{x}, \mathbf{z}) = p_{\beta_0}(\mathbf{x}|\mathbf{z})p_{\alpha, \beta_{>0}}(\mathbf{z})$, and one for data density,

i.e., $q_{\omega}(\mathbf{x}, \mathbf{z}) = p_{\text{data}}(\mathbf{x})q_{\omega}(\mathbf{z}|\mathbf{x})$. We propose joint learning through KL minimization, denoting the objective to be $L(\theta, \omega)$, i.e.,

$$\min_{\theta} \min_{\omega} L(\theta, \omega) = \min_{\theta} \min_{\omega} D_{\text{KL}}(q_{\omega}(\mathbf{x}, \mathbf{z}) || p_{\theta}(\mathbf{x}, \mathbf{z})) \quad (15)$$

Learning generation model β_0 : For learning β_0 , we can compute the gradient as

$$-\nabla_{\beta_0} L(\theta, \omega) = \mathbb{E}_{p_{\text{data}}(\mathbf{x})} \mathbb{E}_{q_{\omega}(\mathbf{z}|\mathbf{x})} [\nabla_{\beta_0} \log p_{\beta_0}(\mathbf{x}|\mathbf{z})] \quad (16)$$

Learning prior model $\alpha, \beta_{>0}$: For learning α_i , the gradient is computed as

$$\begin{aligned} -\nabla_{\alpha_i} L(\theta, \omega) &= \mathbb{E}_{p_{\text{data}}(\mathbf{x})} \mathbb{E}_{q_{\omega}(\mathbf{z}|\mathbf{x})} [\nabla_{\alpha_i} f_{\alpha_i}(\mathbf{z}_i)] \\ &\quad - \mathbb{E}_{p_{\alpha, \beta_{>0}}(\mathbf{z})} [\nabla_{\alpha_i} f_{\alpha_i}(\mathbf{z}_i)] \end{aligned} \quad (17)$$

For learning $\beta_{>0}$, we compute the gradient as

$$\begin{aligned} -\nabla_{\beta_i} L(\theta, \omega) &= \mathbb{E}_{p_{\text{data}}(\mathbf{x})} \mathbb{E}_{q_{\omega}(\mathbf{z}|\mathbf{x})} [\nabla_{\beta_i} \log p_{\beta_i}(\mathbf{z}_i|\mathbf{z}_{i+1})] \\ &\quad - \mathbb{E}_{p_{\alpha, \beta_{>0}}(\mathbf{z})} [\nabla_{\beta_i} \log p_{\beta_i}(\mathbf{z}_i|\mathbf{z}_{i+1})] \end{aligned} \quad (18)$$

Learning inference model ω : For learning ω_i , the gradient is

$$\begin{aligned} -\nabla_{\omega_i} L(\theta, \omega) &= \nabla_{\omega_i} \mathbb{E}_{p_{\text{data}}(\mathbf{x})} [\mathbb{E}_{q_{\omega}(\mathbf{z}|\mathbf{x})} \log p_{\beta_0}(\mathbf{x}|\mathbf{z})] \\ &\quad - D_{\text{KL}}(q_{\omega}(\mathbf{z}|\mathbf{x}) || p_{\beta_{>0}}(\mathbf{z})) \\ &\quad + \mathbb{E}_{q_{\omega}(\mathbf{z}|\mathbf{x})} \left[\sum_{i=1}^L f_{\alpha_i}(\mathbf{z}_i) \right] \end{aligned} \quad (19)$$

Divergence Perturbation. The KL joint minimization (Eqn.15) can be viewed as a surrogate of the MLE objective with the KL perturbation term,

$$\begin{aligned} D_{\text{KL}}(q_{\omega}(\mathbf{x}, \mathbf{z}) || p_{\theta}(\mathbf{x}, \mathbf{z})) \\ = D_{\text{KL}}(p_{\text{data}}(\mathbf{x}) || p_{\theta}(\mathbf{x})) + D_{\text{KL}}(q_{\omega}(\mathbf{z}|\mathbf{x}) || p_{\theta}(\mathbf{z}|\mathbf{x})) \end{aligned}$$

where the perturbation term $D_{\text{KL}}(q_{\omega}(\mathbf{z}|\mathbf{x}) || p_{\theta}(\mathbf{z}|\mathbf{x}))$ measures the KL-divergence between inference distribution and generator posterior. The inference model is learned to directly match the posterior distribution of the generator without expensive posterior sampling. In fact, such KL minimization in the joint space is closely related to evidence lower bound (ELBO) with the joint EBM as the prior model.

4. Related Work

Hierarchical VAEs: Variational auto-encoder (VAE) [23] proposes variational learning by introducing an approximation of the true intractable posterior, which allows a tractable bound on log-likelihood to be formed. But the non-hierarchical structure can be limited in model expressivity and fails to capture different levels of abstraction. Hierarchical VAEs (HVAEs) [5, 28, 35, 40] consist of multiple

layers of latent variables on top of the generator as a prior model, which can be used for learning complex data distribution and hierarchical representations. However, such models still focus on layer-wise relations while ignoring the intra-layer contextual relations at each layer.

Energy-based models: The energy-based models receive attention for being expressive and powerful in capturing contextual data regularities. The majority of existing works focus on the pixel space [7, 8, 10, 14, 15, 43–45]. Learning such EBMs can be done using MLE, where MCMC sampling is typically required in each learning iteration which can be computationally expensive upon data space. Instead, [33] proposes to build EBM on latent space where the energy function is considered as a correction of the non-informative Gaussian prior. The low dimensionality of the latent space makes EBM effective in capturing data regularities and can alleviate the burden of MCMC sampling.

Generator models with informative prior: For generator models, the assumed Gaussian or uniform prior distribution can be non-informative and less expressive. To address this problem, recent works [2, 6, 11, 33, 38, 42] propose to learn generator models with an informative prior, where RAE [11] constructs priors using rejection sampling, and Two-stage VAE [6] propose to train an extra model for simple prior at the second stage to match the aggregated posterior distribution, while LEBM [33] and NCP-VAE [2] instead learn EBMs on latent space to improve the expressivity of generator models.

5. Experiments

To demonstrate the proposed method, we present extensive experiments, including (i) image synthesis, (ii) hierarchical representations, and (iii) analysis of latent space, as well as various ablation studies based on the proposed EBM prior. The parameter complexity is discussed in Sec.5.5.

5.1. Image Synthesis

Generator models with informative prior. We first evaluate the generation performance of the proposed joint model. If the model is well-trained, the multi-layer EBM prior model should render an expressive prior distribution leading to realistic synthesis. We benchmark our model against other generator models that assume standard Gaussian prior, such as VAE [23], Alternating Back-propagation (ABP) [13], Ladder VAE (LVAE) [35], and Short-run Inference (SRI) [32], as well as other generator models using informative prior, such as RAE [11], Two-stages VAE (2s-VAE) [6], NCP-VAE [2], and LEBM [33], where LEBM builds EBM for single layer latent variables, while ours contains a multi-layer structure.

Our project page is available at <https://jcui1224.github.io/hierarchical-joint-ebm-proj>.

Model	SVHN	CelebA-64	CIFAR-10
VAE [23]	46.78	65.75	106.37
LVAE (L=5) [35]	39.26	53.40	-
ABP [13]	49.71	51.50	-
SRI (L=5) [32]	35.32	47.95	-
RAE [11]	42.02	40.95	74.16
2s-VAE [6]	42.81	44.40	72.90
NCP-VAE [2]	33.23	42.07	78.06
LEBM [33]	29.44	37.87	70.15
Ours (L=2)	26.81	33.60	66.32

Table 1. FID(↓) for our model and baselines on SVHN, CelebA (64 x 64), and CIFAR-10.

We train our model on SVHN [29], CIFAR-10 [24] and CelebA-64 [27] and use Fréchet Inception Distance (FID) [17] to quantitatively evaluate the generation quality. To make fair comparisons, we follow the standard protocol as in [33] and use the same generation model with convolutional structures. We use Langevin posterior sampling for the training, and the generation model is jointly learned (the result for variational learning is shown in Ablation Studies). The comparisons are shown in Tab.1, where the superior generation performance indicates the effectiveness of our model in learning a more expressive prior.

Toward deep hierarchical models. We then consider the modern deep hierarchical structures as our multi-layer generator and explore the potential of the joint EBM prior for better generation. We adopt the two-stage training [2, 43] where the deep multi-layer generator $p_{\beta>0}(\mathbf{z})$ and inference model $q_{\omega}(\mathbf{z}|\mathbf{x})$ are trained in the first stage by maximizing the ELBO as in VAEs, with the pre-trained models, our joint EBM prior model can then be learned in the second stage where the posterior samples are directly obtained from the pre-trained inference model $q_{\omega}(\mathbf{z}|\mathbf{x})$ and prior samples can be obtained via Langevin sampling with change of variable on the generator $p_{\beta>0}(\mathbf{z})$.

We consider NVAE [40], a modern hierarchical VAE, for the first stage training, and we train our joint EBM prior in the second stage. For prior sampling in the second stage training, we employ similar *reparametrized sampling* scheme as in [43] via provided code¹ in order to better traverse the deep hierarchical latent space with different scales. We examine our model on CIFAR-10, CelebA-HQ-256 [19], and LSUN-Church-64 [46]. The qualitative results for CelebA-HQ-256 and LSUN-Church-64 are shown in Fig.2 and Fig.3. For CelebA-HQ-256, we synthesize with adjusted batch-normalization as used in [2, 43]. We also visualize the Langevin transition on CIFAR-10 in Fig.4 where the quality of synthesis improves as the Langevin progresses.

¹<https://github.com/NVlabs/VAEBM>



Figure 2. Generated samples on CelebA-HQ-256 . FID = 9.89.



Figure 3. Generated samples on LSUN-Church-64. FID = 8.38.




Figure 4. Langevin transitions on CIFAR-10. FID = 11.34.

Method	IS	FID
NVAE* [40]	5.30	37.73
Ours	8.99	11.34
NCP-VAE [2]	-	24.08
VAEBM [43]	8.43	12.19
Other EBMs		
IGEBM [8]	6.78	38.2
ImprovedCD [7]	7.85	25.1
Divergence Triangle [14]	-	30.10
Adv-EBM [45]	9.10	13.21
Other Likelihood Models		
GLOW [22]	3.92	48.9
PixelCNN [41]	4.60	65.93
GANs+Score-based Models		
BigGAN [3]	9.22	14.73
StyleGANv2 w/o ADA [20]	8.99	9.9
NCSN [36]	8.87	25.32
DDPM [18]	9.46	3.17

Table 2. IS(\uparrow) and FID(\downarrow) for our model and baselines on CIFAR-10. Model* indicates our backbone model.

The quantitative results are shown in Tab.2 and Tab.3. We consider the baseline models, including NCP-VAE [2] and VAEBM [43], which also recruit NVAE as their backbone model, and other powerful deep generative models, such as GANs [3, 20], score-based models [18, 36] and EBMs [7, 8, 14, 45] on data space. Compared to NVAE backbone model, our joint EBM prior model can significantly improve the fidelity of generated samples while only accounting for negligible overhead (see Parameter Efficiency in Sec.5.5). In comparison with other powerful deep generative models, we also achieve competitive generation performance.

Model	CelebA-HQ-256	LSUN-Church-64
NVAE* [40]	30.25	38.13
Ours	9.89	8.38
NCP-VAE [2]	24.79	-
VAEBM [43]	20.38	13.51
Adv-EBM [45]	17.31	10.84
GLOW [22]	68.93	59.35
PGGAN [19]	8.03	6.42

Table 3. FID(\downarrow) for our model and baselines on CelebA-HQ-256 and LSUN-Church-64. Model* indicates backbone model.

5.2. Hierarchical Representations

Hierarchical sampling. To examine our model in learning hierarchical representation, we employ hierarchical sampling to illustrate the learned representation at different layers. In particular, we first sample one group of latent vectors from EBM prior and hold them as fixed constants, then we randomly sample multiple groups of latent vectors to replace the fixed latent vectors at different layers. This allows us to visualize the variation in representation across layers.

We apply our joint EBM prior to BIVA [28] on CelebA-64. For training, we reuse the two-stage training scheme, where we recruit BIVA² for the first stage training and train our EBM prior model in the second stage by using the *reparametrized sampling* [43] method (similar as deep hierarchical models in Sec.5.1). We show the results of hierarchical sampling in Fig.6 and observe that BIVA presents minor changes at the bottom and middle layers, while the proposed joint EBM prior model can show variations of different levels. Note that it is a challenging task for conditional hierarchical models [48], on which the improvement thus suggests that the proposed method is capable of learning hierarchical representations for multi-layer generator models.

²<https://github.com/vlievin/biva-pytorch>




Figure 5. Histograms of density of $L_{\text{EBM}}^{>k}$ with AUROC(\uparrow) and AUPRC(\uparrow) for CIFAR-10 (in) / SVHN (out).




Figure 6. Hierarchical sampling for Gaussian prior model (bottom) and EBM prior model (top). From left panel to right panel, latent vectors are sampled from the bottom layers to the top layers. Detailed sampling method can be referred in Figure 3 in [48].

Out-of-distribution detection. Next, we conduct out-of-distribution (OOD) detection to further evaluate the hierarchical representations. Typically, low-level representations (e.g., edges, corners) can be shared across data which in turn leads to high-confidence reconstructions for OOD examples, while high-level semantic ones have fewer correlations across different data and shall be more discriminative for OOD detection. Inspired by [16], we consider an unnormalized log-posterior as the decision function for EBM prior model, which is defined as

$$L_{\text{EBM}}^{>k} = \mathbb{E}_{\mathbf{z}_{>k} \sim q_{\omega}(\mathbf{z}|\mathbf{x}), \mathbf{z}_{\leq k} \sim p_{\beta_{>0}, \alpha}(\mathbf{z})} [\log p_{\beta_0}(\mathbf{x}|\mathbf{z}) + \log p_{\beta_{>0}}(\mathbf{z}) + \sum_{i=1}^L f_{\alpha_i}(\mathbf{z}_i)] \quad (20)$$

where latent codes above the k -th layer are inferred from inference model and kept fixed, and those below the k -th layer are sampled from EBM prior via the *reparametrized sampling*³ [43] with fixed inferred latent codes. With $k = 0$, all layers of latent vectors are inferred from $q_{\omega}(\mathbf{z}|\mathbf{x})$. With a higher value of k , less inferred low-level representations are used, which should render better performance in OOD detection. In addition, we can also compute a subtraction

between $L_{\text{EBM}}^{>0}$ and $L_{\text{EBM}}^{>k}$ as a surrogate of the likelihood-ratio which is shown to be effective for OOD detection [16]. We compute the subtraction as

$$LLR_{\text{EBM}}^{>k} = L_{\text{EBM}}^{>0} - L_{\text{EBM}}^{>k} \quad (21)$$

We follow standard protocols and apply our EBM prior model with BIVA on CIFAR-10 and use SVHN as OOD data for testing. In Fig.5, we show the density of in-distribution and OOD data by computing the unnormalized log-posterior with increased k , and we use AUROC, AUPRC to quantitatively evaluate the performance. It can be seen that as k increases, relatively lower log-likelihoods are assigned to OOD data, which in turn renders better detection performance (higher AUROC and AUPRC). More importantly, we observe that the backbone model BIVA achieves the best detection performance of 0.885 for AUROC, while our models achieve 0.927 with the adapted decision function. This further verifies that the hierarchical representations can be learned within our multi-layer structure.

5.3. Analysis of Latent Space

Long-run langevin transition. In this section, we examine the energy landscape of our joint EBM prior model. If the EBM is well-learned, the energy prior should naturally render local modes of the energy function, and traversing these local modes should present realistic synthesized examples and steady-state energy scores. Existing EBMs typically have oversaturated images via long-run Langevin dynamics as observed in [30]. Training an EBM that learns steady-state energy scores over realistic images can be useful but challenging.

We train our model on CelebA-64 using Langevin dynamic for 40 steps. We then run 100 and 2500 Langevin steps to examine the learned energy landscape. We show the synthesis and corresponding energy profile in Fig.7. It can be seen that generated examples become sharper for the first 40 steps as it starts from the referenced distribution $p_{\beta_{>0}}(\mathbf{z})$ toward the learned energy prior $p_{\beta_{>0}, \alpha}(\mathbf{z})$, and the energy fluctuates around some constant. For long-run 2500 steps, it is worth noting that our EBM prior model delivers diverse and realistic synthesis, and it does not exhibit the oversaturated phenomenon. This suggests that the learned EBM could mix well between different local modes of the learned energy prior.

³<https://github.com/NVlabs/VAEBM>




Figure 7. Trajectory in data space and energy profile in Langevin transition. **Top:** Langevin transition with 100 steps. **Bottom:** Langevin transition with 2500 steps.

Anomaly Detection. We further evaluate how our joint EBM prior model could benefit the anomaly detection (AD) task. Different from OOD detection, AD requires one class (e.g., one-digit class from MNIST) of data to be held out as anomaly for training, and both normal (e.g., other nine-digit classes from MNIST) and anomalous data are used for testing.

The proposed prior model is built on the joint of all layers of latent variables. If it is well learned, the posterior $q_{\omega}(z|x)$ could form a discriminative joint latent space that has separated probability densities for normal and anomalous data. We use un-normalized log-posterior $L_{EBM}^{>0}$ as our decision function and train our model on MNIST with each class held out as an anomalous class. We consider the baseline models that also adopt an inferential mechanism, such as VAE [23], MEG [25], BiGAN- σ [47], OT-SRI [1], and LEBM [33] which assumes single-layer latent space and is closely related to our method. Tab.4 shows the results of AUPRC scores averaged over the last 10 epochs to account for the variance. To make fair comparisons, we follow the protocols in [1, 25, 33, 47].

Heldout Digit	1	4	5	7	9
VAE [23]	0.063	0.337	0.325	0.148	0.104
MEG [25]	0.281 \pm 0.035	0.401 \pm 0.061	0.402 \pm 0.062	0.290 \pm 0.040	0.342 \pm 0.034
BiGAN- σ [47]	0.287 \pm 0.023	0.443 \pm 0.029	0.514 \pm 0.029	0.347 \pm 0.017	0.307 \pm 0.028
OT-SRI [1]	0.353 \pm 0.021	0.770 \pm 0.024	0.726 \pm 0.030	0.550 \pm 0.013	0.555 \pm 0.023
LEBM [33]	0.336 \pm 0.008	0.630 \pm 0.017	0.619 \pm 0.013	0.463 \pm 0.009	0.413 \pm 0.010
Ours	0.470 \pm 0.009	0.941 \pm 0.001	0.964 \pm 0.003	0.815 \pm 0.004	0.796 \pm 0.004

Table 4. AUPRC scores for unsupervised anomaly detection.

5.4. Ablation Studies.

MCMC sampling vs. Inference model. Two posterior sampling schemes using MCMC and inference model are compared in Tab.5 in terms of FID and wall-clock training time (per-iteration). The MCMC posterior sampling renders better FID as it is more accurate in inference [13, 32], but it can be computationally heavy. While the inference model

is efficient in learning but can be less accurate. For deep hierarchical structures, the variational learning with inference model is preferred due to its efficiency.

MCMC / Inf	SVHN	CelebA-64	CIFAR-10
FID	26.81 / 28.60	33.60 / 36.12	66.32 / 68.45
Time(s)	0.478 / 0.232	0.920 / 0.246	0.568 / 0.256

Table 5. FID and training time for MCMC posterior sampling and variational learning.

Langevin steps. We explore the different number of Langevin steps in prior sampling for training on CIFAR-10. The results of FID and corresponding training time are shown in Tab.6. We observe that the Langevin step k increasing from 10 to 40 can improve the generation quality, while for steps more than 40, it only has minor impacts on the improvement but with increased training overhead. We thus report the result of $k = 40$ in Tab.1.

steps k	$k = 10$	$k = 20$	$k = 40$	$k = 80$	$k = 100$
FID	69.42	67.58	66.32	66.03	65.86
Time(s)	0.312	0.480	0.568	0.741	0.837

Table 6. FID and training time for increasing MCMC steps in prior sampling.

5.5. Parameter Efficiency

It is crucial to analyze the parameter complexity when comparing the generation performance. In Tab.1, we build our model with two layers of latent variables on top of the generator used in [33]. The additional layer accounts for only 1% overhead in total parameter complexity compared to LEBM [33]. For deep hierarchical models, we apply our joint EBM prior model on latent space which brings minimum overhead. The parameter complexity of the backbone NVAE and our EBM model is shown in Tab.7.

NVAE / EBM	CIFAR-10	CelebA-HQ-256	LSUN-Church-64
FID	39.73 / 11.34	30.25 / 9.89	38.13 / 8.38
Parameters	257M / 9M (3%)	375M / 18M (4%)	65M / 5M (7%)

Table 7. FID and parameter complexity for backbone model and EBM.

6. Conclusion

we propose a joint EBM prior for multi-layer generator models, which can effectively capture the intra-layer relations at each layer and jointly correct the latent variables from all layers. We present a joint training scheme via MLE and further develop a variational learning scheme for efficient inference. Our comprehensive experiments demonstrate the effectiveness of the proposed method.

References

- [1] Dongsheng An, Jianwen Xie, and Ping Li. Learning deep latent variable models by short-run mcmc inference with optimal transport correction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15415–15424, June 2021. 8
- [2] Jyoti Aneja, Alex Schwing, Jan Kautz, and Arash Vahdat. A contrastive learning approach for training variational autoencoder priors. *Advances in neural information processing systems*, 34:480–493, 2021. 2, 3, 5, 6
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 6
- [4] Yuri Burda, Roger B. Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. 2
- [5] Rewon Child. Very deep vaes generalize autoregressive models and can outperform them on images. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. 1, 4
- [6] Jian Dai and David Wipf. Diagnosing and enhancing vae models. *arXiv preprint arXiv:1903.05789*, 2019. 5
- [7] Yilun Du, Shuang Li, Joshua Tenenbaum, and Igor Mordatch. Improved contrastive divergence training of energy based models. *arXiv preprint arXiv:2012.01316*, 2020. 2, 5, 6
- [8] Yilun Du and Igor Mordatch. Implicit generation and generalization in energy-based models. *arXiv preprint arXiv:1903.08689*, 2019. 2, 5, 6
- [9] Conor Durkan and Charlie Nash. Autoregressive energy machines. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 1735–1744. PMLR, 2019. 3
- [10] Ruiqi Gao, Yang Song, Ben Poole, Ying Nian Wu, and Diederik P Kingma. Learning energy-based models by diffusion recovery likelihood. *arXiv preprint arXiv:2012.08125*, 2020. 2, 5
- [11] Partha Ghosh, Mehdi SM Sajjadi, Antonio Vergari, Michael Black, and Bernhard Schölkopf. From variational to deterministic autoencoders. *arXiv preprint arXiv:1903.12436*, 2019. 5
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 2
- [13] Tian Han, Yang Lu, Song-Chun Zhu, and Ying Nian Wu. Alternating back-propagation for generator network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017. 2, 5, 8
- [14] Tian Han, Erik Nijkamp, Xiaolin Fang, Mitch Hill, Song-Chun Zhu, and Ying Nian Wu. Divergence triangle for joint training of generator model, energy-based model, and inferential model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 5, 6
- [15] Tian Han, Erik Nijkamp, Linqi Zhou, Bo Pang, Song-Chun Zhu, and Ying Nian Wu. Joint training of variational auto-encoder and latent energy-based model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7978–7987, 2020. 5
- [16] Jakob D Drachmann Havtorn, Jes Frellesen, Soren Hauberg, and Lars Maaløe. Hierarchical vaes know what they don’t know. In *International Conference on Machine Learning*, pages 4117–4128. PMLR, 2021. 7
- [17] 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 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 5
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 6
- [19] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 5, 6
- [20] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems*, 33:12104–12114, 2020. 6
- [21] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 1
- [22] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018. 6
- [23] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2, 4, 5, 8
- [24] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 5
- [25] Rithesh Kumar, Sherjil Ozair, Anirudh Goyal, Aaron Courville, and Yoshua Bengio. Maximum entropy generators for energy-based models. *arXiv preprint arXiv:1901.08508*, 2019. 8
- [26] Don S Lemons and Anthony Gythiel. Paul langevin’s 1908 paper “on the theory of brownian motion”[“sur la théorie du mouvement brownien,” *cr acad. sci.(paris)* 146, 530–533 (1908)]. *American Journal of Physics*, 65(11):1079–1081, 1997. 4
- [27] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. *CoRR*, abs/1411.7766, 2014. 5
- [28] Lars Maaløe, Marco Fraccaro, Valentin Liévin, and Ole Winther. Biva: A very deep hierarchy of latent variables for

- generative modeling. *Advances in neural information processing systems*, 32, 2019. 1, 4, 6
- [29] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bis-sacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011. 5
- [30] Erik Nijkamp, Mitch Hill, Tian Han, Song-Chun Zhu, and Ying Nian Wu. On the anatomy of mcmc-based maximum likelihood learning of energy-based models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5272–5280, 2020. 7
- [31] Erik Nijkamp, Mitch Hill, Song-Chun Zhu, and Ying Nian Wu. Learning non-convergent non-persistent short-run mcmc toward energy-based model. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [32] Erik Nijkamp, Bo Pang, Tian Han, Linqi Zhou, Song-Chun Zhu, and Ying Nian Wu. Learning multi-layer latent variable model via variational optimization of short run mcmc for approximate inference. In *European Conference on Computer Vision*, pages 361–378. Springer, 2020. 1, 2, 5, 8
- [33] Bo Pang, Tian Han, Erik Nijkamp, Song-Chun Zhu, and Ying Nian Wu. Learning latent space energy-based prior model. *Advances in Neural Information Processing Systems*, 33:21994–22008, 2020. 1, 2, 5, 8
- [34] Masaki Saito, Shunta Saito, Masanori Koyama, and So-suke Kobayashi. Train sparsely, generate densely: Memory-efficient unsupervised training of high-resolution tempo-ral gan. *International Journal of Computer Vision*, 128(10):2586–2606, 2020. 1
- [35] Casper Kaae Sønderby, Tapani Raiko, Lars Maaløe, Søren Kaae Sønderby, and Ole Winther. Ladder varia-tional autoencoders. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Infor-mation Processing Systems*, volume 29. Curran Associates, Inc., 2016. 1, 2, 4, 5
- [36] Yang Song and Stefano Ermon. Generative modeling by esti-mating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019. 6
- [37] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Ab-hishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equa-tions. *arXiv preprint arXiv:2011.13456*, 2020. 1
- [38] Jakub Tomczak and Max Welling. Vae with a vampprior. In *International Conference on Artificial Intelligence and Statistics*, pages 1214–1223. PMLR, 2018. 5
- [39] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1
- [40] Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical vari-ational autoencoder. *Advances in Neural Information Pro-cessing Systems*, 33:19667–19679, 2020. 1, 4, 5, 6
- [41] Aaron Van Den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *Interna-tional conference on machine learning*, pages 1747–1756. PMLR, 2016. 6
- [42] Zhisheng Xiao and Tian Han. Adaptive multi-stage den-sity ratio estimation for learning latent space energy-based model. *arXiv preprint arXiv:2209.08739*, 2022. 5
- [43] Zhisheng Xiao, Karsten Kreis, Jan Kautz, and Arash Vah-dat. Vaebm: A symbiosis between variational autoencoders and energy-based models. *arXiv preprint arXiv:2010.00654*, 2020. 2, 5, 6, 7
- [44] Jianwen Xie, Yaxuan Zhu, Jun Li, and Ping Li. A tale of two flows: Cooperative learning of langevin flow and nor-malizing flow toward energy-based model. In *International Conference on Learning Representations*, 2022. 5
- [45] Xuwang Yin, Shiyong Li, and Gustavo K Rohde. Analyzing and improving generative adversarial training for generative modeling and out-of-distribution detection. *arXiv preprint arXiv:2012.06568*, 2020. 2, 5, 6
- [46] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianx-iong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 5
- [47] Houssam Zenati, Chuan Sheng Foo, Bruno Lecouat, Gau-rav Manek, and Vijay Ramaseshan Chandrasekhar. Ef-ficient gan-based anomaly detection. *arXiv preprint arXiv:1802.06222*, 2018. 8
- [48] Shengjia Zhao, Jiaming Song, and Stefano Ermon. Learning hierarchical features from generative models. *arXiv preprint arXiv:1702.08396*, 2017. 1, 6, 7