

# Denoising Diffusion Variational Inference: Diffusion Models as Expressive Variational Posteriors

Wasu Top Piriyaikulij<sup>\*1</sup>, Yingheng Wang<sup>\*1</sup>, Volodymyr Kuleshov<sup>1,2</sup>

<sup>1</sup>Department of Computer Science, Cornell University

<sup>2</sup>The Jacobs Technion-Cornell Institute, Cornell Tech  
{wp237, yw2349, kuleshov}@cornell.edu

## Abstract

We propose denoising diffusion variational inference (DDVI), a black-box variational inference algorithm for latent variable models which relies on diffusion models as flexible approximate posteriors. Specifically, our method introduces an expressive class of diffusion-based variational posteriors that perform iterative refinement in latent space; we train these posteriors with a novel regularized evidence lower bound (ELBO) on the marginal likelihood inspired by the wake-sleep algorithm. Our method is easy to implement (it fits a regularized extension of the ELBO), is compatible with black-box variational inference, and outperforms alternative classes of approximate posteriors based on normalizing flows or adversarial networks. We find that DDVI improves inference and learning in deep latent variable models across common benchmarks as well as on a motivating task in biology—inferring latent ancestry from human genomes—where it outperforms strong baselines on the Thousand Genomes dataset.

## 1 Introduction

We are interested in amortized black-box variational inference problems of the form

$$\log p_{\theta}(\mathbf{x}) \geq \max_{\phi} \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})} [\log p_{\theta}(\mathbf{x}, \mathbf{z}) - \log q_{\phi}(\mathbf{z}|\mathbf{x})] \\ := \max_{\phi} \text{ELBO}(\mathbf{x}, \theta, \phi), \quad (1)$$

in which we approximate the marginal likelihood  $\log p_{\theta}(\mathbf{x})$  of a latent variable model  $p_{\theta}(\mathbf{x}, \mathbf{z})$  with an evidence lower bound  $\text{ELBO}(\mathbf{x}, \theta, \phi)$  that is a function of an approximate variational posterior  $q_{\phi}(\mathbf{z}|\mathbf{x})$ . We assume that  $p_{\theta}$  factorizes as  $p_{\theta}(\mathbf{x}|\mathbf{z})p_{\theta}(\mathbf{z})$  and admits efficient sampling: examples of such  $p_{\theta}$  include Bayesian networks, topic models (Blei, Ng, and Jordan 2003), variational autoencoders (VAEs), and broad classes of  $p_{\theta}$  defined via modern probabilistic programming frameworks (Gordon et al. 2014).

Maximizing  $\text{ELBO}(\mathbf{x}, \theta, \phi)$  over  $\phi$  yields a variational posterior  $q_{\phi}(\mathbf{z}|\mathbf{x})$  that approximates  $p_{\theta}(\mathbf{z}|\mathbf{x})$  as well as a tight bound on  $\log p_{\theta}(\mathbf{x})$  that serves as a learning objective for  $p_{\theta}$ . The approximation gap  $\log p_{\theta}(\mathbf{x}) - \max_{\phi} \text{ELBO}(\mathbf{x}, \theta, \phi)$  equals precisely

$\min_{\phi} \text{KL}(q_{\phi}(\mathbf{z}|\mathbf{x})||p_{\theta}(\mathbf{z}|\mathbf{x}))$ , which motivates the design of expressive classes of posteriors  $q_{\phi}(\mathbf{z}|\mathbf{x})$  that reduce this gap. Recent efforts leverage modern generative models—including normalizing flows (Rezende and Mohamed 2015; Kingma et al. 2016) and generative adversarial networks (Goodfellow et al. 2014; Makhzani et al. 2015)—as expressive model families for  $q_{\phi}$  that tighten the ELBO.

This work seeks to further improve variational inference via expressive posteriors based on diffusion models (Ho, Jain, and Abbeel 2020; Song et al. 2020). Diffusion methods have become the de-facto standard for high-quality image synthesis (Rombach et al. 2022; Gokaslan et al. 2024). Here, we use diffusion in latent space to parameterize  $q_{\phi}(\mathbf{z}|\mathbf{x})$ . We train this distribution with a denoising diffusion-like objective that does not involve adversarial training (Makhzani et al. 2015) or constrained invertible normalizing flow architectures (Kingma et al. 2016). Samples from  $q_{\phi}(\mathbf{z}|\mathbf{x})$  are obtained via iterative refinement of  $\mathbf{z}$ , starting from a Gaussian distribution, and gradually forming one that is multi-modal and complex.

Our work expands upon existing diffusion-based approximate inference methods (Berner, Richter, and Ullrich 2022; Zhang and Chen 2021; Vargas, Grathwohl, and Doucet 2023; Zhang et al. 2023; Richter, Berner, and Liu 2023; Sendera et al. 2024; Akhound-Sadegh et al. 2024) that focus on the task of drawing samples from unnormalized distributions  $\tilde{p}(\mathbf{z})$  and estimating the partition function  $Z = \int_{\mathbf{z}} \tilde{p}(\mathbf{z}) d\mathbf{z}$ . While these methods are applicable in our setting—we set the unnormalized  $\tilde{p}(\mathbf{z})$  to  $p_{\theta}(\mathbf{x}, \mathbf{z})$  such that  $Z = p_{\theta}(\mathbf{x})$ —they do not make use of characteristics of  $p_{\theta}$  that are common in many types of models (VAEs, Bayes networks, etc.), namely the factorization  $p_{\theta}(\mathbf{x}|\mathbf{z})p_{\theta}(\mathbf{z})$  and efficient sampling. We find that leveraging these properties yields simpler algorithms that avoid backpropagating through a sampling process, and that are fast enough to perform learning in addition to inference.

Specifically, we propose denoising diffusion variational inference (DDVI), an approximate inference algorithm defined by a class of approximate posterior distribution based on diffusion and a learning objective inspired by the wake-sleep algorithm (Hinton et al. 1995) that implements regularized variational inference. We also derive extensions of our method to semi-supervised learning and clustering.

Our method is easy to implement (it fits a regularized ex-

<sup>\*</sup>These authors contributed equally.

tension of the ELBO), is compatible with black-box variational inference, and outperforms alternative classes of approximate posteriors based on normalizing flows or adversarial networks.

We evaluate DDVI on synthetic benchmarks and on a real problem in biological data analysis—inferring human ancestry from genetic data. Our method outperforms strong baselines on the Thousand Genomes dataset (Siva 2008) and learns a low-dimensional latent space that preserves biologically meaningful structure (Hagverdi, Buettner, and Theis 2015).

**Contributions.** In summary, this work introduces denoising diffusion variational inference, an approximate inference algorithm defined by two components: a class of approximate posteriors  $q(\mathbf{z}|\mathbf{x})$  parameterized by diffusion, and a lower bound on the marginal likelihood inspired by wake-sleep. We complement DDVI with extensions to semi-supervised learning and clustering. Our method is especially suited for probabilistic programming, representation learning, and dimensionality reduction, where it outperforms alternative methods based on normalizing flows and adversarial training.

## 2 Background

**Deep Latent Variable Models** Latent variable models (LVMs)  $p_\theta(\mathbf{x}, \mathbf{z})$  are usually fit by optimizing the evidence lower bound (ELBO)

$$\log p_\theta(\mathbf{x}) \geq \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \text{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z})),$$

which serves as a tractable surrogate for the marginal log-likelihood (MLL). The gap between the MLL and the ELBO equals precisely  $\text{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z}|\mathbf{x}))$ —thus, a more expressive  $q_\phi(\mathbf{z}|\mathbf{x})$  may better fit the true posterior and induce a tighter ELBO (Kingma and Welling 2013).

Expressive variational posteriors can be formed by choosing more expressive model families—including auxiliary variable methods (Maaløe et al. 2016), MCMC-based methods (Salimans, Kingma, and Welling 2015), normalizing flows (Rezende and Mohamed 2015)—or improved learning objectives—e.g., adversarial or sample-based losses (Makhzani et al. 2015; Zhao, Song, and Ermon 2017; Si, Bishop, and Kuleshov 2022; Si et al. 2023).

The wake-sleep algorithm (Hinton et al. 1995) optimizes an alternative objective

$$\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \text{KL}(p_\theta(\mathbf{z}|\mathbf{x})||q_\phi(\mathbf{z}|\mathbf{x})),$$

in which the KL divergence term is reversed. The learning procedure for wake-sleep involves alternating between “wake” phases where the recognition model is updated and “sleep” phases where the generative model is refined.

**Denoising Diffusion Models** A diffusion model is defined via a user-specified noising process  $q$  that maps data  $\mathbf{x}_0$  into a sequence of  $T$  variables  $\mathbf{y}_{1:T} = \mathbf{y}_1, \dots, \mathbf{y}_T$  that represent increasing levels of corruption to  $\mathbf{x}_0$ . We obtain  $\mathbf{y}_{1:T}$  by applying a Markov chain  $q(\mathbf{y}_{1:T}|\mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{y}_t|\mathbf{y}_{t-1})$ , where we define  $\mathbf{y}_0 = \mathbf{x}_0$  for convenience. When  $\mathbf{x}_0$  is a continuous vector, a standard choice of transition kernel is

$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{y}_t; \sqrt{\alpha_t}\mathbf{y}_{t-1}, \sqrt{1 - \alpha_t}\mathbf{I})$ , which is a Gaussian centered around a copy of  $\mathbf{y}_{t-1}$  to which we added noise following a schedule  $0 < \alpha_1 < \alpha_2 < \dots < \alpha_T = 1$ .

A diffusion model can then be represented as a latent variable distribution  $p(\mathbf{x}_0, \mathbf{y}_{1:T})$  that factorizes as  $p(\mathbf{x}_0, \mathbf{y}_{1:T}) = p(\mathbf{y}_T) \prod_{t=0}^{T-1} p_\theta(\mathbf{y}_t | \mathbf{y}_{t+1})$  (again using  $\mathbf{y}_0$  as shorthand for  $\mathbf{x}_0$ ). This model seeks to approximate the reverse of the forward diffusion  $q$  and map noise  $\mathbf{y}_T$  into data  $\mathbf{x}_0$ .

The true reverse of the process  $q$  cannot be expressed in closed form; as such, we parameterize  $p_\theta$  with  $\theta$  trained by maximizing the ELBO:

$$\log p_\theta(\mathbf{x}_0) \geq \mathbb{E}_q \left[ \log p_\theta(\mathbf{x}_0|\mathbf{x}_1) - \sum_{t=2}^T \text{KL}(q_t||p_t) \right] - \text{KL}(q(\mathbf{x}_T|\mathbf{x}_0)||p(\mathbf{x}_T)) \quad (2)$$

where  $q_t, p_t$  denote the distributions  $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$  and  $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$ , respectively.

## 3 Variational Inference With Denoising Diffusion Models

We introduce *denoising diffusion variational inference (DDVI)*, which improves variational inference with diffusion-based techniques.

The goal of DDVI is to fit a latent variable model  $p_\theta(\mathbf{x}, \mathbf{z})$ . We assume that  $p_\theta$  factorizes as  $p_\theta(\mathbf{x}|\mathbf{z})p_\theta(\mathbf{z})$  and admits efficient sampling: examples of such  $p_\theta$  include Bayesian networks and variational autoencoders (VAEs) (Kingma and Welling 2013).

Our approach is comprised of three components:

1. A **modeling family** of approximate posteriors  $q_\phi(\mathbf{z}|\mathbf{x})$  based on diffusion;
2. A **learning objective** formed by a regularized ELBO;
3. An **optimization algorithm** inspired by wake-sleep.

The  $q_\phi(\mathbf{z}|\mathbf{x})$  iteratively refines latents  $\mathbf{z}$ , starting from a Gaussian distribution. The learning objective trains  $q_\phi(\mathbf{z}|\mathbf{x})$  to reverse a user-specified forward diffusion process.

### 3.1 Modeling Family: Diffusion-Based Posteriors

DDVI performs variational inference using a family of approximate posteriors  $q_\phi(\mathbf{z}|\mathbf{x}) = \int_{\mathbf{y}} q_\phi(\mathbf{z}|\mathbf{y}, \mathbf{x})q_\phi(\mathbf{y}|\mathbf{x})d\mathbf{y}$ , which themselves contain latent variables  $\mathbf{y} \in \mathcal{Y}$ . The models  $q_\phi(\mathbf{z}|\mathbf{y}, \mathbf{x}), q_\phi(\mathbf{y}|\mathbf{x})$  must have tractable densities and support gradient-based optimization over  $\phi$ .

We choose the latent  $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_T)$  to be a vector of  $T$  variables that represent progressively simplified versions of  $\mathbf{z}$ , with  $\mathbf{y}_T$  corresponding to a simple distribution (e.g., a Gaussian). The model  $q_\phi(\mathbf{y}, \mathbf{z}|\mathbf{x}) = q_\phi(\mathbf{z}|\mathbf{y}_1, \mathbf{x}) \prod_{t=1}^{T-1} q_\phi(\mathbf{y}_t|\mathbf{y}_{t+1}, \mathbf{x})$  transforms  $\mathbf{y}_T$  into  $\mathbf{z}$  via iterative refinement. To sample from  $q_\phi$ , we first sample  $\mathbf{y}_T$ —this is an easier task since we can define  $\mathbf{y}_T$  to have a simple (e.g., Gaussian) distribution—and then by sampling from the denoising model  $q_\phi(\mathbf{z}|\mathbf{y}_1, \mathbf{x}) \prod_{t=1}^{T-1} q_\phi(\mathbf{y}_t|\mathbf{y}_{t+1}, \mathbf{x})$ .

We define the relationship between  $\mathbf{y}$  and  $\mathbf{z}$  via a forward diffusion process  $r(\mathbf{y}|\mathbf{z}, \mathbf{x}) = r(\mathbf{y}_{1:T}|\mathbf{z}, \mathbf{x}) =$

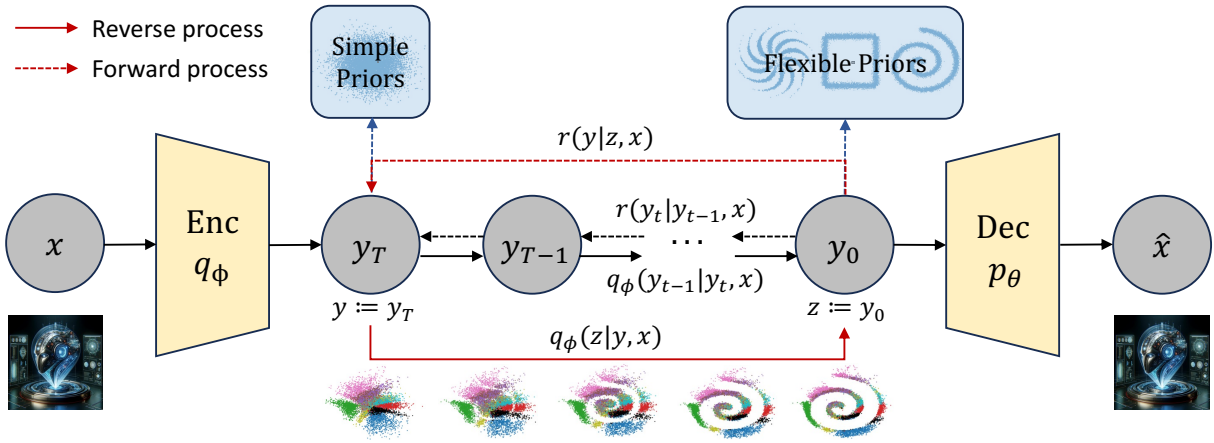


Figure 1: Denoising diffusion variational inference in a VAE. Between the encoder and decoder, we have a diffusion model to map a simple distribution into a complex distribution over latents.

$r(\mathbf{y}_1|\mathbf{z}, \mathbf{x}) \prod_{t=1}^{T-1} r(\mathbf{y}_{t+1}|\mathbf{y}_t, \mathbf{x})$ , which transforms  $\mathbf{z}$ —the latent whose intractable posterior we seek to approximate—into  $\mathbf{y}_T$ , whose posterior is easier to model (possibly conditioned on  $\mathbf{x}$ ). Examples of  $r$  include Gaussian forward diffusion processes and discrete noising processes (Austin et al. 2021). The model  $q_\phi$  is trained to approximately reverse this forward diffusion process.

### 3.2 Learning Objective: A Markovian ELBO

The standard approach to fit auxiliary-variable generative models (Maaløe et al. 2016) is to apply the ELBO twice:

$$\log p_\theta(\mathbf{x}) \geq \log p_\theta(\mathbf{x}) - \text{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z}|\mathbf{x})) \quad (3)$$

$$\begin{aligned} &\geq \log p_\theta(\mathbf{x}) - \text{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z}|\mathbf{x})) \\ &\quad - \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\text{KL}(q_\phi(\mathbf{y}|\mathbf{x}, \mathbf{z})||r(\mathbf{y}|\mathbf{x}, \mathbf{z}))] \end{aligned} \quad (4)$$

$$\begin{aligned} &= \mathbb{E}_{q_\phi(\mathbf{y}, \mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] \\ &\quad - \text{KL}(q_\phi(\mathbf{y}, \mathbf{z}|\mathbf{x})||r(\mathbf{y}|\mathbf{x}, \mathbf{z})p(\mathbf{z})) \end{aligned} \quad (5)$$

In Equation (3), we applied the ELBO over  $\mathbf{z}$ , and in Equation (4) we applied the ELBO over the latent  $\mathbf{y}$  of  $q$  (see Appendix J for the derivation). Notice that the gap between the ELBO and  $\log p_\theta(\mathbf{x})$  is  $\text{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p_\theta(\mathbf{z}|\mathbf{x})) + \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\text{KL}(q_\phi(\mathbf{y}|\mathbf{x}, \mathbf{z})||r(\mathbf{y}|\mathbf{x}, \mathbf{z}))]$ . Thus, if we correctly match  $q$  and  $r$ , we will achieve a tight bound.

**Analyzing the ELBO** Optimizing Equation (5) requires tractably dealing with the prior regularization term  $\mathcal{L}_{\text{reg}}(\mathbf{x}, \theta, \phi) := -\text{KL}(q_\phi(\mathbf{y}, \mathbf{z}|\mathbf{x})||r(\mathbf{y}|\mathbf{x}, \mathbf{z})p(\mathbf{z}))$ , which we equivalently rewrite as:

$$\mathcal{L}_{\text{reg}} = \mathbb{E}_{q_\phi(\mathbf{y}, \mathbf{z}|\mathbf{x})}[\log(r(\mathbf{y}|\mathbf{x}, \mathbf{z})p(\mathbf{z}))] + H(q). \quad (6)$$

We can expand the first term by leveraging the Markov structure of  $r, q$  to rewrite  $\mathcal{L}_{\text{reg}}$  as the likelihood of samples from the reverse diffusion process  $q$  under the forward process  $r$ .

$$\mathcal{L}_{\text{reg}} = \sum_{t=1}^T \mathbb{E}_q[\log(r(\mathbf{y}_t|\mathbf{y}_{t-1}, \mathbf{x}))] + \mathbb{E}_q[\log p(\mathbf{z})] + H(q),$$

where  $\mathbf{y}_0 := \mathbf{z}$ . We refer to optimizing the Markovian ELBO as **unregularized DDVI**, the first instance of our method.

The noise process  $r$  defines prior regularization terms for each  $\mathbf{y}_t$ . This provides extra supervision for learning  $q$  in the form of trajectories from latents  $\mathbf{y}_T$  to  $\mathbf{y}_1$ ; this extra supervision helps  $q$  fit complex non-Gaussian posteriors.

The term  $H(q) = -\sum_{t=1}^{T+1} \mathbb{E}_q[\log q_\phi(\mathbf{y}_{t-1}|\mathbf{y}_t, \mathbf{x})]$  denotes the entropy. For example, when each term  $q_\phi(\mathbf{y}_{t-1}|\mathbf{y}_t, \mathbf{x})$  is Gaussian, it is computed as

$$H(q) = \sum_{t=1}^{T+1} \mathbb{E}_q \left[ \frac{d}{2} (1 + \log(2\pi)) + \frac{1}{2} \log |\Sigma_\phi(\mathbf{y}_t, \mathbf{x})| \right]$$

where  $d$  is the dimension of  $\mathbf{y}$  and we use the notation  $\mathbf{y}_{T+1} = \mathbf{x}$ . It is also common to leave the variance  $\Sigma_\phi$  fixed, in which case  $H(q)$  is a constant.

### 3.3 Refining the Objective: A Regularized ELBO

Notice that optimizing  $\mathcal{L}_{\text{reg}}$  involves sampling from the approximate reverse process  $q_\phi(\mathbf{y}, \mathbf{z}|\mathbf{x})$  to match the true reverse process  $r(\mathbf{y}|\mathbf{z}, \mathbf{x})$ : this is the opposite of diffusion training, where we would sample from  $r$  to fit  $q$ . This type of on-policy learning of  $q$  has been studied in the context of approximate inference (Zhang and Chen 2021); however, it requires backpropagating through  $T$  samples, which may hamper training, and it optimizes a mode-covering divergence that may struggle to fit complex  $p(\mathbf{z})$ .

**Adding Wake-Sleep Regularization to the ELBO** We propose to further improve the ELBO via off-policy diffusion-like training. Our new objective is the ELBO in Equation (5) augmented with a regularizer  $\mathcal{L}_{\text{sleep}}(\phi)$ .

$$\begin{aligned} \log p_\theta(\mathbf{x}) \geq &\underbrace{\mathbb{E}_{q_\phi(\mathbf{y}, \mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})]}_{\text{wake / recons. term } \mathcal{L}_{\text{rec}}(\mathbf{x}, \theta, \phi)} \\ &\underbrace{-\text{KL}(q_\phi(\mathbf{y}, \mathbf{z}|\mathbf{x})||r(\mathbf{y}|\mathbf{x}, \mathbf{z})p(\mathbf{z}))}_{\text{prior regularization term } \mathcal{L}_{\text{reg}}(\mathbf{x}, \theta, \phi)} \\ &\underbrace{-\mathbb{E}_{p_\theta(\mathbf{x})}[\text{KL}(p_\theta(\mathbf{z}|\mathbf{x})||q_\phi(\mathbf{z}|\mathbf{x}))]}_{\text{sleep term } \mathcal{L}_{\text{sleep}}(\phi)} \end{aligned} \quad (7)$$

The optimization of the regularizer  $\mathcal{L}_{\text{sleep}}(\phi)$  is similar to the sleep phase of wake-sleep, and closely resembles diffusion model training (see below). As in wake-sleep,  $\mathcal{L}_{\text{sleep}}(\phi)$  is optimized over  $\phi$  only, the  $\mathbf{x}$  are sampled from  $p$ .

**From Wake-Sleep to Diffusion Regularization** Computing  $\mathcal{L}_{\text{sleep}}(\phi)$  still involves intractable distributions  $p_{\theta}(\mathbf{z}|\mathbf{x}), q_{\phi}(\mathbf{z}|\mathbf{x})$ . To optimize  $\mathcal{L}_{\text{sleep}}(\phi)$ , we introduce another lower bound  $\mathcal{L}_{\text{diff}}(\phi)$ , which we call the denoising diffusion loss (for reasons that will become apparent shortly):

$$\begin{aligned}\mathcal{L}_{\text{sleep}}(\phi) &= -\mathbb{E}_{p_{\theta}(\mathbf{x})}[\text{KL}(p_{\theta}(\mathbf{z}|\mathbf{x})||q_{\phi}(\mathbf{z}|\mathbf{x}))] \quad (8) \\ &= \mathbb{E}_{p_{\theta}(\mathbf{x}, \mathbf{z})}[\log q_{\phi}(\mathbf{z}|\mathbf{x})] + \bar{H}(p_{\theta}) \\ &\geq \mathbb{E}_{p_{\theta}(\mathbf{x}, \mathbf{z})} \left[ \mathbb{E}_r \left[ \log \frac{q_{\phi}(\mathbf{y}, \mathbf{z}|\mathbf{x})}{r(\mathbf{y}|\mathbf{z}, \mathbf{x})} \right] \right] + \bar{H}(p_{\theta}) \\ &= \mathcal{L}_{\text{diff}}(\phi)\end{aligned}$$

In Equation (8), we applied the ELBO with  $r(\mathbf{y}|\mathbf{z}, \mathbf{x})$  playing the role of the variational posterior over the latent  $\mathbf{y}$  in  $q_{\phi}$ ;  $\bar{H}(p_{\theta})$  is the expected conditional entropy of  $p_{\theta}(\mathbf{z}|\mathbf{x})$ , a constant that does not depend on  $\phi$ .

We can further simplify  $\mathcal{L}_{\text{diff}}(\phi)$  by leveraging the Markov structure of the forward and reverse processes  $r, q$ . Recall that each  $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_T)$  can be a vector of  $T$  latents, which we also denote as  $\mathbf{y}_{1:T}$ , and that  $r(\mathbf{y}_{1:T}|\mathbf{z}, \mathbf{x}) = \prod_{t=1}^T r(\mathbf{y}_t|\mathbf{y}_{t-1}, \mathbf{x})$ , where  $\mathbf{y}_0 = \mathbf{z}$  and also  $q_{\phi}(\mathbf{y}, \mathbf{z}|\mathbf{x}) = q_{\phi}(\mathbf{y}_{0:T}|\mathbf{x}) = q_{\phi}(\mathbf{y}_T|\mathbf{x}) \prod_{t=1}^T q_{\phi}(\mathbf{y}_{t-1}|\mathbf{y}_t, \mathbf{x})$ .

We may use the Markov structure in  $q, r$  to rewrite  $\mathcal{L}_{\text{diff}}(\phi)$  as a sum of  $T$  terms, one per Markov step. The derivation is identical to that used to obtain the ELBO of a diffusion model (Sohl-Dickstein et al. 2015), and yields an expression of the same form:

$$\begin{aligned}\mathcal{L}_{\text{diff}}(\phi) &= \mathbb{E}_r \left[ \log q_{\phi}(\mathbf{z}|\mathbf{y}_1, \mathbf{x}) - \sum_{t=2}^T \text{KL}(r_t||q_t) \right] \quad (9) \\ &\quad - \text{KL}(r(\mathbf{y}_T|\mathbf{z})||q_{\phi}(\mathbf{y}|\mathbf{x})).\end{aligned}$$

where  $r_t, q_t$  denote the distributions  $r(\mathbf{y}_{t-1}|\mathbf{y}_t, \mathbf{y}_0, \mathbf{x})$  and  $q_{\phi}(\mathbf{y}_{t-1}|\mathbf{y}_t, \mathbf{x})$  (see Appendix K for the derivation).

**Regularized DDVI Objective** We define the full DDVI objective  $\mathcal{L}_{\text{ddvi}}$  to be the sum of the aforementioned terms:

$$\mathcal{L}_{\text{ddvi}}(\mathbf{x}, \theta, \phi) = \mathcal{L}_{\text{rec}}(\mathbf{x}, \theta, \phi) + \mathcal{L}_{\text{reg}}(\mathbf{x}, \theta, \phi) + \mathcal{L}_{\text{diff}}(\phi)$$

Terms  $\mathcal{L}_{\text{reg}}$  and  $\mathcal{L}_{\text{diff}}$  may be weighted by hyper-parameters  $\beta_{\text{reg}}, \beta_{\text{diff}} > 0$ , as in the  $\beta$ -VAE framework. In our experiments,  $\beta_{\text{reg}} = \beta_{\text{diff}} = 1$  unless otherwise specified. Note that since  $\mathcal{L}_{\text{diff}} \leq \mathcal{L}_{\text{sleep}} \leq 0$ ,  $\mathcal{L}(\mathbf{x}, \theta, \phi)$  is a valid lower bound on  $\log p_{\theta}(\mathbf{x})$  that is tight when  $q_{\phi}(\mathbf{z}|\mathbf{x}) = p_{\theta}(\mathbf{z}|\mathbf{x})$ .

### 3.4 Optimization: Extending Wake-Sleep

We may optimize  $\mathcal{L}_{\text{ddvi}}(\mathbf{x}, \theta, \phi)$  using gradient descent by alternating between ELBO optimization and taking sleep steps (see Appendix A for full details):

1. Sample  $\mathbf{x}$  from data, take gradient step on  $\theta, \phi$  optimizing  $\mathcal{L}_{\text{rec}}(\mathbf{x}, \theta, \phi) + \mathcal{L}_{\text{reg}}(\mathbf{x}, \theta, \phi)$  (the ‘‘wake’’ step);
2. Sample  $\mathbf{z}, \mathbf{x}$  from  $p_{\theta}$  and take a gradient step on  $\phi$  optimizing  $\mathcal{L}_{\text{diff}}(\phi)$  (the ‘‘sleep’’ step).

Again, terms may be weighted by  $\beta_{\text{reg}}, \beta_{\text{diff}} > 0$ . Note that by maximizing  $\mathcal{L}_{\text{diff}}(\phi)$ , we fit  $q_{\phi}(\mathbf{z}|\mathbf{x})$  to  $p_{\theta}(\mathbf{z}|\mathbf{x})$  via the forward KL divergence; similarly, by optimizing  $\mathcal{L}_{\text{rec}} + \mathcal{L}_{\text{reg}}$  (the ELBO), we fit  $q_{\phi}(\mathbf{z}|\mathbf{x})$  to  $p_{\theta}(\mathbf{z}|\mathbf{x})$  via the reverse KL divergence. Thus, optimizing  $\mathcal{L}(\mathbf{x}, \theta, \phi)$  encourages  $q_{\phi}(\mathbf{z}|\mathbf{x})$  to approximate  $p_{\theta}(\mathbf{z}|\mathbf{x})$ , and when the two are equal, the bound  $\mathcal{L}_{\text{ddvi}}$  on  $\log p_{\theta}(\mathbf{x})$  is tight.

**Simplified Wake-Sleep** We also consider a light-weight algorithm, in which  $r(\mathbf{y}|\mathbf{z})$  and  $q_{\phi}(\mathbf{z}|\mathbf{y})$  do not depend on  $\mathbf{x}$ . This scenario admits the following optimization procedure:

1. Sample  $\mathbf{x}$  from data and compute gradient on  $\theta, \phi$  optimizing  $\mathcal{L}_{\text{rec}}(\mathbf{x}, \theta, \phi) + \mathcal{L}_{\text{reg}}(\mathbf{x}, \theta, \phi)$ .
2. Sample  $\mathbf{z}$  from  $p(\mathbf{z})$  and compute gradient on  $\phi$  optimizing  $\mathcal{L}_{\text{diff}}(\phi)$ ; take step on weighted sum of both gradients.

In this case,  $\mathcal{L}_{\text{diff}}$  requires only sampling from  $p(\mathbf{z})$ , and the entire loss  $\mathcal{L}_{\text{ddvi}}$  can be optimized end-to-end using gradient descent. This algorithm is simpler (there is no separate sleep phase); however,  $q_{\phi}(\mathbf{z}|\mathbf{x})$  may not perfectly approximate  $p_{\theta}(\mathbf{z}|\mathbf{x})$  when  $r(\mathbf{y}|\mathbf{z})$  and  $q_{\phi}(\mathbf{z}|\mathbf{y})$  do not depend on  $\mathbf{x}$ , hence  $\mathcal{L}$  may no longer be a tight bound.

**Practical Considerations** A common type of noising process compatible with this bound when  $\mathbf{z}$  is continuous is Gaussian diffusion, where we define  $r(\mathbf{y}_t|\mathbf{y}_{t-1}) = \mathcal{N}(\mathbf{y}_t; \sqrt{1 - \alpha_t}\mathbf{y}_{t-1}, \alpha_t\mathbf{I})$  for a suitable schedule  $(\alpha_t)_{t=1}^T$ . We then adopt the parameterization  $q_{\phi}(\mathbf{y}_{t-1}|\mathbf{y}_t, \mathbf{x}) = \mathcal{N}(\mathbf{y}_{t-1}; \mu_{\phi}(\mathbf{y}_t, \mathbf{x}, t), \Sigma_{\phi}(\mathbf{y}_t, \mathbf{x}, t))$ . It is then common to parameterize  $q_{\phi}$  with a noise prediction network  $\epsilon_{\phi}$  (Ho, Jain, and Abbeel 2020); the sum of KL divergences can be approximated by  $\mathbb{E}_{t, \epsilon_t \sim r(\mathbf{y}_0, t)} \|\epsilon_t - \epsilon_{\phi}(\sqrt{\alpha_t}\mathbf{y}_0 + \sqrt{1 - \alpha_t}\epsilon_t, \mathbf{x}, t)\|^2$ . Other extensions include discrete denoising processes (Austin et al. 2021; Sahoo et al. 2024; Schiff et al. 2024). In the wake-sleep setting, we know both endpoints  $\mathbf{y}_T \sim q(\cdot|\mathbf{x})$  and  $\mathbf{y}_0 = \mathbf{z}$  of the diffusion process, opening the possibility for applying optimal transport techniques (Cuturi 2013; De Bortoli et al. 2021).

## 4 Extensions

### 4.1 Semi-Supervised Learning

Following Makhzani et al. (2015), we extend our algorithm to the semi-supervised learning setting where some data points have labels denoted by  $l$ . We assume the user provides a model of the form  $p_{\theta}(\mathbf{x}, \mathbf{y}, \mathbf{z}, l) = p_{\theta}(\mathbf{x}|\mathbf{z}, l)r(\mathbf{y}|\mathbf{z}, l)p_{\theta}(\mathbf{z}|l)p(l)$ ; we set the variational distributions to  $q_{\phi}(\mathbf{z}|\mathbf{x}, \mathbf{y}, l), q_{\phi}(\mathbf{y}|\mathbf{x}), q_{\phi}(l|\mathbf{x})$ . In this setting, we consider two cases, depending on whether the label is observed (Kingma et al. 2014). We extend Equation (7) to incorporate the label  $l$  corresponding to a data point as follows:

$$\begin{aligned}\mathcal{L}_{\text{semi}} &= \mathbb{E}_{q_{\phi}(\mathbf{y}, \mathbf{z}|\mathbf{x}, l)} [\log p_{\theta}(\mathbf{x}|\mathbf{z}, l)] \quad (10) \\ &\quad - \text{KL}(q_{\phi}(\mathbf{y}, \mathbf{z}|\mathbf{x}, l)||p_{\theta}(\mathbf{y}, \mathbf{z}|l)) \\ &\quad - \mathbb{E}_{p_{\theta}(\mathbf{x})} [\text{KL}(p_{\theta}(\mathbf{z}|\mathbf{x}, l)||q_{\phi}(\mathbf{z}|\mathbf{x}, l))]\end{aligned}$$

When the label  $c$  cannot be observed, we treat it as a latent variable and modify the learning objective  $\mathcal{U}_{\text{semi}} = \sum_c q_{\phi}(l|\mathbf{x})\mathcal{L}_{\text{semi}}(\mathbf{x}, l, \theta, \phi) + \text{KL}(q_{\phi}(l|\mathbf{x})||p(l))$ . Therefore, we can conclude a marginal likelihood on our

dataset as follows:  $\tilde{\mathcal{L}}_{\text{semi}} = \sum_{(\mathbf{x}, l) \in L} \mathcal{L}_{\text{semi}}(\mathbf{x}, l, \boldsymbol{\theta}, \phi) + \sum_{\mathbf{x} \in U} \mathcal{U}_{\text{semi}}(\mathbf{x}, \boldsymbol{\theta}, \phi)$ . where  $L$  and  $U$  are the sets of data with and without labels, respectively.

We also want to guarantee that all model parameters can be learned in all cases, including  $q_{\phi}(l|\mathbf{x})$ , such that this posterior can be applied as a classifier during inference. Thus, we combine the marginal likelihood with a classification loss to form an extended learning objective:  $\tilde{\mathcal{L}}_{\text{semi}, \alpha} = \tilde{\mathcal{L}}_{\text{semi}} + \alpha \cdot \mathbb{E}_{\tilde{p}(\mathbf{x}, l)} [-\log q_{\phi}(l|\mathbf{x})]$

## 4.2 Clustering

We have further extended our algorithm to encompass the clustering paradigm. We propose two distinct strategies. In the first approach, we simply formulate a model in which  $p_{\boldsymbol{\theta}}(\mathbf{z})$  is a mixture of desired priors. The means of these priors are characterized by  $\boldsymbol{\theta}$ . From these means, cluster membership, denoted as  $\mathbf{c}$  can be deduced. This approach requires no alteration to the existing learning objective.

Alternatively, the second method retains the original prior, but introduces an additional latent cluster variable  $\mathbf{c}$  where  $\sum_i c_i = 1$ . Thus, the model can be specified as  $p_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{c}) = p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}, \mathbf{c})r(\mathbf{y}|\mathbf{z})p_{\boldsymbol{\theta}}(\mathbf{z})p(\mathbf{c})$  with  $p(\mathbf{c}) = \text{Dir}(\boldsymbol{\epsilon})$ . Consequently, the variational distributions become  $q_{\phi}(\mathbf{z}|\mathbf{y}, \mathbf{c}, \mathbf{x})$ ,  $q_{\phi}(\mathbf{y}, \mathbf{c}|\mathbf{x})$ . This yields the objective:

$$\begin{aligned} \mathcal{L}_{\text{clus}}(\mathbf{x}) = & \mathbb{E}_{q_{\phi}(\mathbf{y}, \mathbf{z}, \mathbf{c}|\mathbf{x})} [\log p_{\boldsymbol{\theta}}(\mathbf{x}|\mathbf{z}, \mathbf{c})] & (11) \\ & - \text{KL}(q_{\phi}(\mathbf{y}, \mathbf{z}, \mathbf{c}|\mathbf{x}) || p_{\boldsymbol{\theta}}(\mathbf{y}, \mathbf{z}, \mathbf{c})) \\ & - \mathbb{E}_{p_{\boldsymbol{\theta}}(\mathbf{x})} [\text{KL}(p_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}) || q_{\phi}(\mathbf{z}|\mathbf{x}))] \end{aligned}$$

Expectations over small numbers of classes  $\mathbf{c}$  are done analytically; larger  $\mathbf{c}$  require backpropagating through discrete sampling (Jang, Gu, and Poole 2016; Sahoo et al. 2023).

## 5 Experiments

We compare DDVI with Auto-Encoding Variational Bayes (AEVB) (Kingma and Welling 2013), AEVB with inverse autoregressive flow posteriors (AEVB-IAF) (Kingma et al. 2016), Adversarial Auto-Encoding Bayes (AAEB) (Makhzani et al. 2015), and Path Integral Sampler (PIS) (Zhang and Chen 2021) on MNIST (Lecun et al. 1998) and CIFAR-10 (Krizhevsky and Hinton 2009) in unsupervised and semi-supervised learning settings, and also on the Thousand Genomes dataset (Siva 2008). We also compare with Hierarchical Auto-Encoding Variational Bayes (H-AEVB) (Ranganath, Tran, and Blei 2016; Vahdat and Kautz 2020) in unsupervised setting. We discuss the computational costs of all methods in Appendix D. The priors, model architecture, and training details can also be found in Appendix H. All results below are reported with 95% confidence interval using 3 different seeds.

### 5.1 Unsupervised learning

We start with synthetic experiments that are aimed at benchmarking the expressivity of diffusion-based posteriors and their ability to improve fitting  $p$ , a distribution with a complex structured prior, like one might find in probabilistic programming, scientific analysis, or other applications. We fit a

model  $p_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{z})$  on the MNIST and CIFAR-10 datasets with three priors  $p(\mathbf{z})$ : pinwheel, swiss roll, and square and report our results in Table 1 and 7. The model distribution  $p_{\boldsymbol{\theta}}$  is instantiated by a deep Gaussian latent variable model (DGLVM) with multi-layer perceptrons (MLPs) on MNIST and convolutional neural networks (CNNs) on CIFAR-10 (see the details of model architecture in Appendix G).

Our first set of metrics (ELBO and MMD) seeks to evaluate the learned generative model  $p_{\boldsymbol{\theta}}$  is good. In the ELBO calculation, we average the reconstruction loss across image pixels. We use MMD to measure sample quality: we generate images with the trained model and calculate MMD between the generated images and test images using a mixture of Gaussian kernel with sigma equal to [2, 5, 10, 20, 40, 80]. We only report MMD for MNIST, since CIFAR-10 generated samples are very low-quality for all methods because the latent dimension is 2.

Our last metric seeks to directly evaluate the expressivity of the posterior. We measure latent negative log-likelihood (Latent NLL) by fitting a kernel density estimator (KDE) on the latents produced by the model with test data as input and compute the log-likelihood of the latents sampled from the prior under the fitted KDE.

From Table 1 and Table 7 in Appendix, we see our method DDVI achieve best ELBO in all but one scenario, in which it still performs competitively. We also see strong results in Latent NLL and k-nearest neighbors classification accuracy of the latents (Acc). in many scenarios, except for swiss roll where AAEB does well. We present visualizations on MNIST using the baselines and our method in Figure 2.

### 5.2 Semi-supervised Learning

We also evaluate the performance of our method and the baselines under semi-supervised learning setting where some labels are observed (1,000 for MNIST and 10,000 for CIFAR-10) and the partitions of the priors are known.

For this setting, we evaluate ELBO, Latent NLL, and Acc. We choose classification accuracy since classification is a common downstream task for semi-supervised learning. We use the same set of priors and baselines. Details on how we partition each prior into  $p_{\boldsymbol{\theta}}(\mathbf{z}|\mathbf{x}, l)$  can be founded in Appendix F. The partitions defined for our priors are local parts of the priors. We note that unlike unsupervised learning, we use the simplified sleep term in our objective for this setting (see Appendix B for details), since  $q_{\phi}$  already gets extra information from  $l$  here.

The results are shown in Table 2 and Table 8 in Appendix. DDVI mostly outperforms the baselines across different priors and metrics, especially on CIFAR-10 where DDVI is best across the board. For MNIST, DDVI always achieves the best ELBO, and it also performs competitively with other baselines in classification accuracy. We also show the visualizations of the latents in Figure 3 where DDVI matches the prior almost perfectly.

### 5.3 Clustering and Visualization for Genotype Analysis

In this section, we report results on a real-world task in genome analysis. Visualizing genotype data reveals patterns

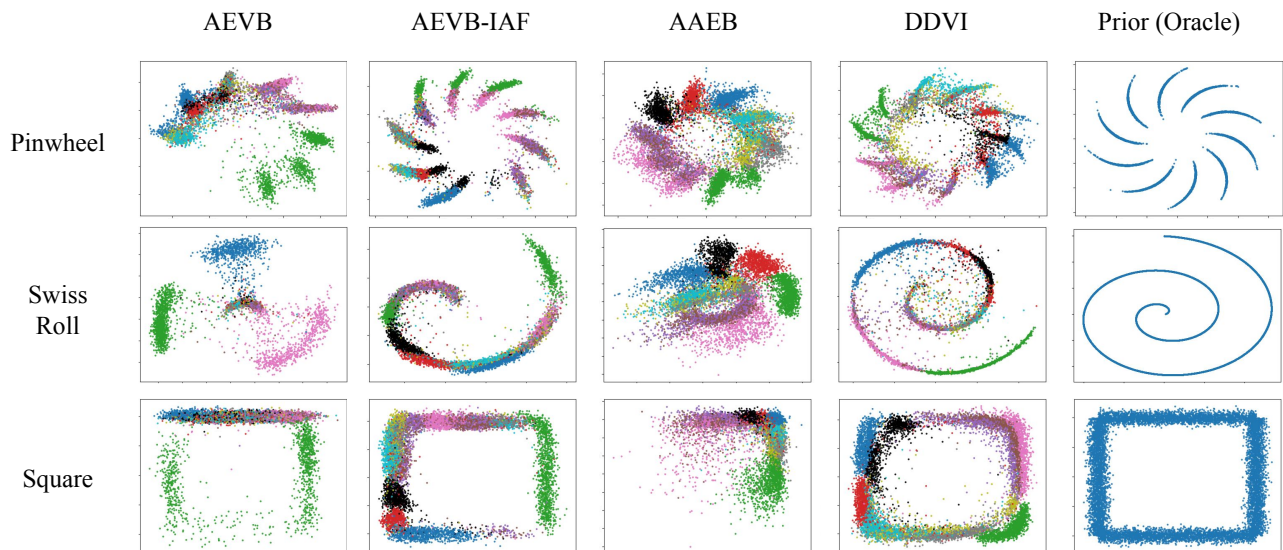


Figure 2: Unsupervised visualization on MNIST using three priors (pinwheel, swiss roll, and square). Each color indicates a class.

Method	Pinwheel			Swiss Roll			Square		
	ELBO	MMD	Latent NLL	ELBO	MMD	Latent NLL	ELBO	MMD	Latent NLL
AEVB	$-12.13 \pm 0.41$	$0.77 \pm 0.04$	$1.68 \pm 0.31$	$-14.80 \pm 0.23$	$0.78 \pm 0.17$	$5.65 \pm 1.58$	$-7.85 \pm 0.29$	$1.10 \pm 0.66$	$2.78 \pm 0.61$
AEVB-IAF	$-4.19 \pm 0.05$	$0.77 \pm 0.00$	$1.64 \pm 0.73$	$-5.10 \pm 0.30$	$0.61 \pm 0.15$	$4.43 \pm 1.09$	$-3.97 \pm 0.22$	$0.75 \pm 0.12$	$1.68 \pm 0.27$
AAEB	N/A	$0.68 \pm 0.02$	$1.54 \pm 0.19$	N/A	<b><math>0.52 \pm 0.03</math></b>	$3.34 \pm 0.16$	N/A	$0.80 \pm 0.02$	$2.46 \pm 0.46$
H-AEBV	$-7.03 \pm 3.13$	$0.74 \pm 0.02$	$2.25 \pm 3.02$	$-7.21 \pm 4.62$	$0.70 \pm 0.22$	$4.04 \pm 4.62$	$-5.71 \pm 3.05$	$0.76 \pm 0.21$	$2.22 \pm 2.03$
PIS	$-7.83 \pm 0.64$	$0.75 \pm 0.14$	$6.50 \pm 1.11$	$-9.83 \pm 0.61$	$0.61 \pm 0.03$	<b><math>2.40 \pm 1.01</math></b>	$-7.06 \pm 0.06$	$0.77 \pm 0.04$	$3.67 \pm 0.08$
DDVI	<b><math>-3.88 \pm 0.96</math></b>	<b><math>0.67 \pm 0.04</math></b>	<b><math>1.27 \pm 0.21</math></b>	<b><math>-5.03 \pm 0.58</math></b>	$0.62 \pm 0.33$	$3.86 \pm 0.17$	<b><math>-3.79 \pm 0.14</math></b>	<b><math>0.66 \pm 0.07</math></b>	<b><math>1.56 \pm 0.09</math></b>

Table 1: Unsupervised learning on MNIST. We report ELBO, MMD between generated images and test images, and latent negative log-likelihood (Latent NLL) with pinwheel, swiss roll, and square priors.

in the latent ancestry of individuals. We instantiate DDVI with a deep Gaussian latent variable model (DGLVM) and compare it against with the three strong clustering baselines using the 1000 Genomes dataset. We also report visualizations from three dimensionality reduction algorithms: PCA, TSNE, and UMAP. For each clustering algorithm, we seek to discover up to 20 clusters. We report quantitative results in terms of cluster purity, cluster completeness, and normalized mutual information (NMI). There is an inherent trade-off between cluster purity completeness. The overall clustering performance can be captured with NMI.

In Table 3, we see that DDVI attains the best performance on cluster purity and NMI. For cluster completeness, VAE and AAE have better means but much larger confidence interval. Furthermore, we visualize our genotype clustering results in latent space, shown in Figure 4, and also report results from classical dimensionality reduction and visualization methods that do not perform clustering (PCA (Wold, Esbensen, and Geladi 1987), t-SNE (Van der Maaten and Hinton 2008), and UMAP (McInnes, Healy, and Melville 2018)). The legend of Figure 4 can be founded at Figure 5 in Appendix.

## 6 Discussion

**Diffusion vs. Normalizing Flows** Our approach is most similar to flow-based approximators (Rezende and Mohamed 2015; Kingma et al. 2016); in fact when  $T \rightarrow \infty$ , our diffusion-based posterior effectively becomes a continuous-time normalizing flow (Song, Meng, and Ermon 2020). However, classical flow-based methods require invertible architectures for each flow layer: this constrains their expressivity and requires backpropagating through potentially a very deep network. Our approach, on the other hand, trains a model (a continuous-time flow when  $T \rightarrow \infty$ ) via a denoising objective (similar to score matching) that does not require invertible architectures and effectively admits an infinite number of layers (with weight sharing). This model is trained not by backpropagating through the ELBO, but rather via an auxiliary diffusion loss term (effectively, a score matching objective).

Despite training with a modified loss, we observe in Section 5 that a diffusion model with an expressive denoising architecture yields an improved ELBO relative to regular flows. Also, our modified loss based on the forward KL divergence reduces posterior collapse (i.e., all modes of the prior are covered well), and thus produces better samples.

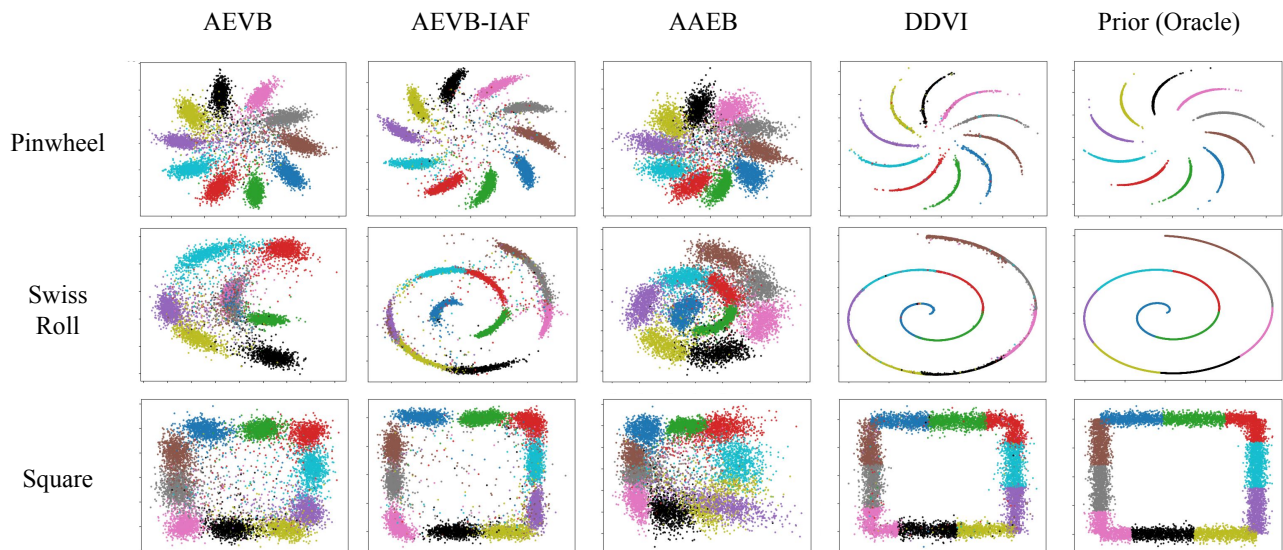


Figure 3: Semi-supervised visualization on MNIST with 1,000 labels using three different priors (pinwheel, swiss roll, and square). Each a indicates one class.

Method	Pinwheel			Swiss Roll			Square		
	ELBO	Acc	Latent NLL	ELBO	Acc	Latent NLL	ELBO	Acc	Latent NLL
AEVB	$-11.15 \pm 0.53$	$0.93 \pm 0.01$	$1.36 \pm 0.03$	$-15.29 \pm 1.33$	$0.68 \pm 0.01$	$4.60 \pm 0.23$	$-10.26 \pm 0.25$	$0.86 \pm 0.01$	$1.68 \pm 0.02$
AEVB-IAF	$-2.10 \pm 0.26$	<b><math>0.95 \pm 0.00</math></b>	<b><math>1.06 \pm 0.03</math></b>	$-5.38 \pm 1.78$	$0.90 \pm 0.02$	$2.75 \pm 0.14$	$-2.67 \pm 0.83$	<b><math>0.91 \pm 0.01</math></b>	<b><math>0.90 \pm 0.02</math></b>
AAEB	<i>N/A</i>	$0.89 \pm 0.01$	$1.55 \pm 0.01$	<i>N/A</i>	$0.88 \pm 0.01$	$3.07 \pm 0.05$	<i>N/A</i>	$1.94 \pm 0.38$	$0.76 \pm 0.13$
DDVI	<b><math>-0.24 \pm 0.13</math></b>	<b><math>0.95 \pm 0.00</math></b>	<b><math>1.06 \pm 0.01</math></b>	<b><math>-2.89 \pm 0.33</math></b>	<b><math>0.92 \pm 0.01</math></b>	<b><math>2.09 \pm 0.00</math></b>	<b><math>0.02 \pm 0.09</math></b>	$0.90 \pm 0.01$	$1.49 \pm 0.03$

Table 2: Semi-supervised learning on MNIST (1,000 labels). We report ELBO, accuracy using KNN (K=20) classifier (Acc), and latent negative log-likelihood (Latent NLL) with pinwheel, swiss roll, and square priors.

Method	Cluster Purity	Cluster Completeness	NMI
AEVB	$0.28 \pm 0.02$	<b><math>0.78 \pm 0.16</math></b>	$0.59 \pm 0.08$
AEVB-IAF	$0.29 \pm 0.04$	<b><math>0.73 \pm 0.06</math></b>	$0.55 \pm 0.06$
AAEB	$0.37 \pm 0.06$	<b><math>0.76 \pm 0.11</math></b>	$0.63 \pm 0.02$
DDVI	<b><math>0.45 \pm 0.03</math></b>	<b><math>0.75 \pm 0.05</math></b>	<b><math>0.66 \pm 0.04</math></b>

Table 3: Quantitative genotype clustering results.

**Diffusion vs. Other Generative Models** Variational posteriors based on GANs (Makhzani et al. 2015) also admit expressive architectures and require only sample-based access to the prior  $p(\mathbf{z})$ . Our diffusion-based approach admits a more stable loss, and is potentially more expressive, as it effectively supports an infinite number of layers (with shared parameters when  $T \rightarrow \infty$ ). Unlike GANs, our models also admit explicit likelihoods and allow us to compute the ELBO for model evaluation. Our approach is similar to variational MCMC (Salimans, Kingma, and Welling 2015); however, we train with a better objective augmented with a diffusion loss, and we adopt improved architectures with shared weights across layers.

**Diffusion for Approximate Inference** Existing diffusion-based approximate inference methods (Berner, Richter, and Ullrich 2022; Zhang and Chen 2021; Vargas, Grathwohl,

and Doucet 2023; Zhang et al. 2023; Richter, Berner, and Liu 2023; Sendera et al. 2024; Akhound-Sadegh et al. 2024) focus on the task of drawing samples from unnormalized distributions  $\tilde{p}(\mathbf{z})$  and estimating the partition function  $Z = \int_{\mathbf{z}} \tilde{p}(\mathbf{z}) d\mathbf{z}$ . While these methods are applicable in our setting—we set the unnormalized  $\tilde{p}(\mathbf{z})$  to  $p_{\theta}(\mathbf{x}, \mathbf{z})$  such that  $Z = p_{\theta}(\mathbf{x})$ —they also tackle a more challenging problem (drawing samples from energy-based models) in more general classes of models (arbitrary unnormalized distributions). In contrast, we focus on restricted but still important classes of models (VAEs, Bayes networks, etc.), and we solve more challenging sets of tasks (e.g., maximum-likelihood learning) by using properties of  $p_{\theta}$  (the factorization  $p_{\theta}(\mathbf{x}|\mathbf{z})p_{\theta}(\mathbf{z})$  and efficient sampling from  $p_{\theta}$ ).

Our algorithms are also simpler. For example, diffusion sampling methods require backpropagating through a sampling process to minimize the reverse  $\text{KL}(q_{\phi}||p_{\theta})$ , which poses challenges with optimization and credit assignment. Some methods based on Schrodinger bridges require an iterative optimization process generalizing the sinkhorn algorithm or computationally expensive on-policy or off-policy (Malkin et al. 2022) trajectory-based optimization. In contrast, DDVI optimizes the forward  $\text{KL}(p_{\theta}||q_{\phi})$  using simple gradient-based optimization that directly emulates diffusion-based training.

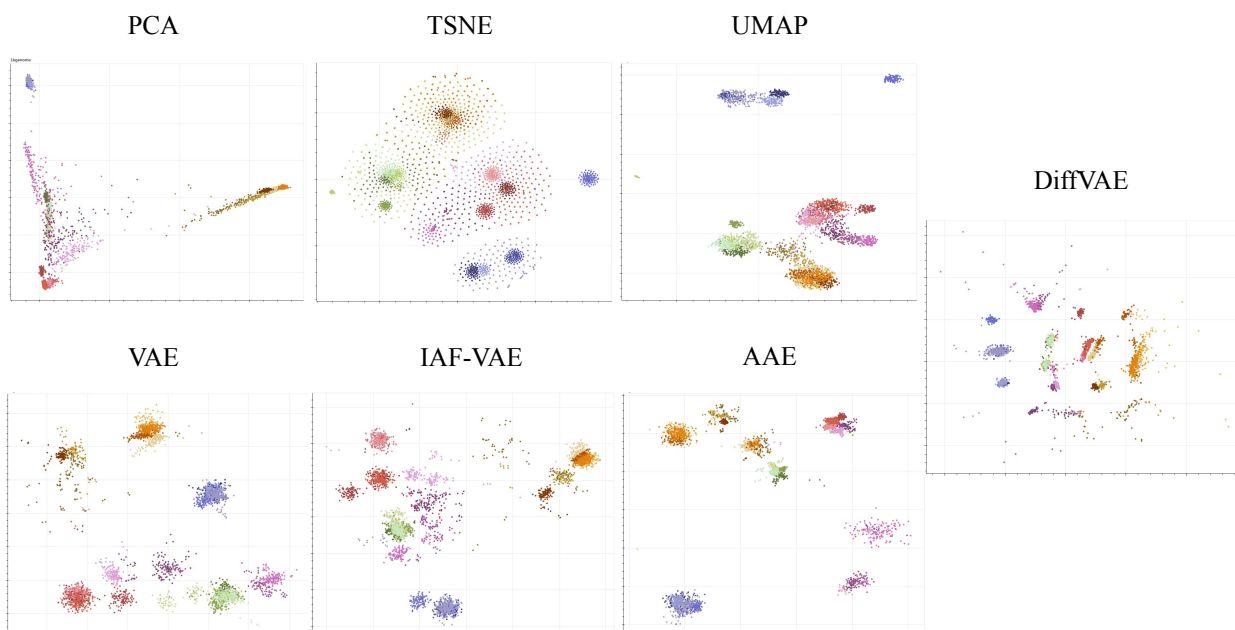


Figure 4: Visualization of genotype clusters. A color represents one ethnicity.

## 7 Related Work

**Latent Diffusion** Vahdat, Kreis, and Kautz (2021); Wehenkel and Louppe (2021); Rombach et al. (2022) perform diffusion in the latent space of a VAE. Their goal is high sample quality, and they introduce into  $p$  hierarchical latents with simple Gaussian priors.

Our goal is different: we seek to fit a  $p$  with structured latents (e.g., in probabilistic programming or in science applications, users introduce prior knowledge via hand-crafted  $p$ ), and we improve variational inference in this structured model by introducing auxiliary latents into  $q$ . Recent work (Preechakul et al. 2022; Zhang, Zhao, and Lin 2022; Wang et al. 2023) has also melded auto-encoders with diffusion models, focusing on semantically meaningful low-dimensional latents in a diffuser  $p$ . Cohen et al. (2022) crafts a diffusion bridge linking a continuously coded vector to a non-informative prior distribution.

**Diffusion for Approximate Inference** Diffusion sampling (Berner, Richter, and Ullrich 2022; Zhang and Chen 2021; Vargas, Grathwohl, and Doucet 2023; Zhang et al. 2023; Richter, Berner, and Liu 2023; Sendera et al. 2024; Akhound-Sadegh et al. 2024) mainly focuses on the task of drawing samples from unnormalized distributions and estimating the partition function. These works draw connections between diffusion (learning the denoising process) and stochastic control (learning the Föllmer drift). Some other works (Zhang et al. 2023; Akhound-Sadegh et al. 2024; Sendera et al. 2024) use continuous generative flow networks (GFlowNets) – deep reinforcement learning algorithms adapted to variational inference that offers stable off-policy training and thus flexible exploration. While this work is applicable to our setting, it does not rely on the structure

of  $p(\mathbf{x}, \mathbf{z})$  available to us, namely tractable sampling in  $p$ .

**Dimensionality Reduction** Latent variable models in general are an attractive alternative to visualization methods like PCA, Sparse PCA, NMF, UMAP, and t-SNE (Wold, Esbensen, and Geladi 1987; Kuleshov 2013; Lee and Seung 2000; McInnes, Healy, and Melville 2018; Van der Maaten and Hinton 2008). Domain-specific knowledge can be injected through the prior, and deep neural networks can be utilized to achieve a more expressive mapping from the data space to the latent space. Nevertheless, downsides of LVMs are that they are more computationally expensive and require careful hyperparameter tuning.

## 8 Conclusion

While this paper focuses on applications of DDVI to dimensionality reduction and visualization, there exist other tasks for the algorithm, e.g., density estimation or sample quality. Accurate variational inference has the potential to improve downstream applications of generative modeling, e.g., decision making (Nguyen and Grover 2022; Deshpande and Kuleshov 2023), meta-learning (Rastogi et al. 2023), or causal effect estimation (Deshpande et al. 2022).

Since our learning objective differs from the ELBO (it adds a regularizer), we anticipate gains on models whose training benefits from regularization, but perhaps not on all models. Also, attaining competitive likelihood estimation requires architecture improvements that are orthogonal to this paper. However, our ability to generate diverse samples and achieve class separation in latent space hints at the method’s potential on these tasks.

## References

- Akhound-Sadegh, T.; Rector-Brooks, J.; Bose, A. J.; Mittal, S.; Lemos, P.; Liu, C.-H.; Sendera, M.; Ravanbakhsh, S.; Gidel, G.; Bengio, Y.; et al. 2024. Iterated Denoising Energy Matching for Sampling from Boltzmann Densities. *arXiv preprint arXiv:2402.06121*.
- Austin, J.; Johnson, D. D.; Ho, J.; Tarlow, D.; and Van Den Berg, R. 2021. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34: 17981–17993.
- Berner, J.; Richter, L.; and Ullrich, K. 2022. An optimal control perspective on diffusion-based generative modeling. *arXiv preprint arXiv:2211.01364*.
- Blei, D. M.; Ng, A. Y.; and Jordan, M. I. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan): 993–1022.
- Cohen, M.; Quispe, G.; Corff, S. L.; Ollion, C.; and Moulines, E. 2022. Diffusion bridges vector quantized Variational AutoEncoders. *arXiv preprint arXiv:2202.04895*.
- Cuturi, M. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26.
- De Bortoli, V.; Thornton, J.; Heng, J.; and Doucet, A. 2021. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34: 17695–17709.
- Deshpande, S.; and Kuleshov, V. 2023. Calibrated Uncertainty Estimation Improves Bayesian Optimization. *arXiv:2112.04620*.
- Deshpande, S.; Wang, K.; Sreenivas, D.; Li, Z.; and Kuleshov, V. 2022. Deep multi-modal structural equations for causal effect estimation with unstructured proxies. *Advances in Neural Information Processing Systems*, 35: 10931–10944.
- Gokaslan, A.; Cooper, A. F.; Collins, J.; Seguin, L.; Jacobson, A.; Patel, M.; Frankle, J.; Stephenson, C.; and Kuleshov, V. 2024. CommonCanvas: Open Diffusion Models Trained on Creative-Commons Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8250–8260.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Gordon, A. D.; Henzinger, T. A.; Nori, A. V.; and Rajamani, S. K. 2014. Probabilistic programming. *Future of software engineering proceedings*, 167–181.
- Haghverdi, L.; Buettner, F.; and Theis, F. J. 2015. Diffusion maps for high-dimensional single-cell analysis of differentiation data. *Bioinformatics*, 31(18): 2989–2998.
- Hinton, G. E.; Dayan, P.; Frey, B. J.; and Neal, R. M. 1995. The “wake-sleep” algorithm for unsupervised neural networks. *Science*, 268(5214): 1158–1161.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Jang, E.; Gu, S.; and Poole, B. 2016. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*.
- Johnson, M. J.; Duvenaud, D.; Wiltchko, A. B.; Datta, S. R.; and Adams, R. P. 2016. Composing Graphical Models with Neural Networks for Structured Representations and Fast Inference. In *Advances in Neural Information Processing Systems (NIPS) 29*. ArXiv:1603.06277 [stat.ML].
- Kingma, D. P.; Mohamed, S.; Jimenez Rezende, D.; and Welling, M. 2014. Semi-supervised learning with deep generative models. *Advances in neural information processing systems*, 27.
- Kingma, D. P.; Salimans, T.; Jozefowicz, R.; Chen, X.; Sutskever, I.; and Welling, M. 2016. Improved variational inference with inverse autoregressive flow. *Advances in neural information processing systems*, 29.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Krizhevsky, A.; and Hinton, G. 2009. Learning multiple layers of features from tiny images. *Toronto, ON, Canada*.
- Kuleshov, V. 2013. Fast algorithms for sparse principal component analysis based on Rayleigh quotient iteration. In *International Conference on Machine Learning*, 1418–1425. PMLR.
- Lecun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.
- Lee, D.; and Seung, H. S. 2000. Algorithms for Non-negative Matrix Factorization. In Leen, T.; Dietterich, T.; and Tresp, V., eds., *Advances in Neural Information Processing Systems*, volume 13. MIT Press.
- Maaløe, L.; Sønderby, C. K.; Sønderby, S. K.; and Winther, O. 2016. Auxiliary deep generative models. In *International conference on machine learning*, 1445–1453. PMLR.
- Makhzani, A.; Shlens, J.; Jaitly, N.; Goodfellow, I.; and Frey, B. 2015. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*.
- Malkin, N.; Jain, M.; Bengio, E.; Sun, C.; and Bengio, Y. 2022. Trajectory balance: Improved credit assignment in GFlowNets. *Advances in Neural Information Processing Systems*, 35: 5955–5967.
- Marsland, S. 2014. *Machine Learning: An Algorithmic Perspective (2nd Edition)*. Chapman and Hall/CRC.
- McInnes, L.; Healy, J.; and Melville, J. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- Nguyen, T.; and Grover, A. 2022. Transformer Neural Processes: Uncertainty-Aware Meta Learning Via Sequence Modeling. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvári, C.; Niu, G.; and Sabato, S., eds., *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, 16569–16594. PMLR.
- Preechakul, K.; Chatthee, N.; Wizadwongsa, S.; and Suwanakorn, S. 2022. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of*

- the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10619–10629.
- Ranganath, R.; Tran, D.; and Blei, D. 2016. Hierarchical variational models. In *International conference on machine learning*, 324–333. PMLR.
- Rastogi, R.; Schiff, Y.; Hacothen, A.; Li, Z.; Lee, I.; Deng, Y.; Sabuncu, M. R.; and Kuleshov, V. 2023. Semi-Parametric Inducing Point Networks and Neural Processes. In *The Eleventh International Conference on Learning Representations*.
- Rezende, D.; and Mohamed, S. 2015. Variational inference with normalizing flows. In *International conference on machine learning*, 1530–1538. PMLR.
- Richter, L.; Berner, J.; and Liu, G.-H. 2023. Improved sampling via learned diffusions. *arXiv preprint arXiv:2307.01198*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Sahoo, S. S.; Arriola, M.; Gokaslan, A.; Marroquin, E. M.; Rush, A. M.; Schiff, Y.; Chiu, J. T.; and Kuleshov, V. 2024. Simple and Effective Masked Diffusion Language Models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Sahoo, S. S.; Paulus, A.; Vlastelica, M.; Musil, V.; Kuleshov, V.; and Martius, G. 2023. Backpropagation through Combinatorial Algorithms: Identity with Projection Works. In *The Eleventh International Conference on Learning Representations*.
- Salimans, T.; Kingma, D. P.; and Welling, M. 2015. Markov Chain Monte Carlo and Variational Inference: Bridging the Gap. In *International conference on machine learning*, 1218–1226. PMLR.
- Schiff, Y.; Sahoo, S. S.; Phung, H.; Wang, G.; Boshar, S.; Dalla-torre, H.; de Almeida, B. P.; Rush, A.; Pierrot, T.; and Kuleshov, V. 2024. Simple Guidance Mechanisms for Discrete Diffusion Models. *arXiv preprint arXiv:2412.10193*.
- Sendera, M.; Kim, M.; Mittal, S.; Lemos, P.; Scimeca, L.; Rector-Brooks, J.; Adam, A.; Bengio, Y.; and Malkin, N. 2024. On diffusion models for amortized inference: Benchmarking and improving stochastic control and sampling. *arXiv preprint arXiv:2402.05098*.
- Si, P.; Bishop, A.; and Kuleshov, V. 2022. Autoregressive Quantile Flows for Predictive Uncertainty Estimation. In *International Conference on Learning Representations*.
- Si, P.; Chen, Z.; Sahoo, S. S.; Schiff, Y.; and Kuleshov, V. 2023. Semi-Autoregressive Energy Flows: Exploring Likelihood-Free Training of Normalizing Flows. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 31732–31753. PMLR.
- Siva, N. 2008. 1000 Genomes project. *Nature biotechnology*, 26(3): 256–257.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, 2256–2265. PMLR.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Vahdat, A.; and Kautz, J. 2020. NVAE: A deep hierarchical variational autoencoder. *Advances in neural information processing systems*, 33: 19667–19679.
- Vahdat, A.; Kreis, K.; and Kautz, J. 2021. Score-based generative modeling in latent space. *Advances in Neural Information Processing Systems*, 34: 11287–11302.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Vargas, F.; Grathwohl, W.; and Doucet, A. 2023. Denoising diffusion samplers. *arXiv preprint arXiv:2302.13834*.
- Wang, Y.; Schiff, Y.; Gokaslan, A.; Pan, W.; Wang, F.; De Sa, C.; and Kuleshov, V. 2023. InfoDiffusion: Representation Learning Using Information Maximizing Diffusion Models. In Krause, A.; Brunskill, E.; Cho, K.; Engelhardt, B.; Sabato, S.; and Scarlett, J., eds., *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, 36336–36354. PMLR.
- Wehenkel, A.; and Louppe, G. 2021. Diffusion priors in variational autoencoders. *arXiv preprint arXiv:2106.15671*.
- Wold, S.; Esbensen, K.; and Geladi, P. 1987. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3): 37–52.
- Zhang, D.; Chen, R. T. Q.; Liu, C.-H.; Courville, A.; and Bengio, Y. 2023. Diffusion generative flow samplers: Improving learning signals through partial trajectory optimization. *arXiv preprint arXiv:2310.02679*.
- Zhang, Q.; and Chen, Y. 2021. Path integral sampler: a stochastic control approach for sampling. *arXiv preprint arXiv:2111.15141*.
- Zhang, Z.; Zhao, Z.; and Lin, Z. 2022. Unsupervised representation learning from pre-trained diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 35: 22117–22130.
- Zhao, S.; Song, J.; and Ermon, S. 2017. Infovae: Information maximizing variational autoencoders. *arXiv preprint arXiv:1706.02262*.