

Causal Diffusion Autoencoders: Toward Counterfactual Generation via Diffusion Probabilistic Models

Aneesh Komanduri^{1,*}, Chen Zhao², Feng Chen³ and Xintao Wu¹

¹University of Arkansas

²Baylor University

³University of Texas at Dallas

Abstract. Diffusion probabilistic models (DPMs) have become the state-of-the-art in high-quality image generation. However, DPMs have an arbitrary noisy latent space with no interpretable or controllable semantics. Although there has been significant research effort to improve image sample quality, there is little work on representation-controlled generation using diffusion models. Specifically, causal modeling and controllable counterfactual generation using DPMs is an underexplored area. In this work, we propose CausalDiffAE, a diffusion-based causal representation learning framework to enable counterfactual generation according to a specified causal model. Our key idea is to use an encoder to extract high-level semantically meaningful causal variables from high-dimensional data and model stochastic variation using reverse diffusion. We propose a causal encoding mechanism that maps high-dimensional data to causally related latent factors and parameterize the causal mechanisms among latent factors using neural networks. To enforce the disentanglement of causal variables, we formulate a variational objective and leverage auxiliary label information in a prior to regularize the latent space. We propose a DDIM-based counterfactual generation procedure subject to do-interventions. Finally, to address the limited label supervision scenario, we also study the application of CausalDiffAE when a part of the training data is unlabeled, which also enables granular control over the strength of interventions in generating counterfactuals during inference. We empirically show that CausalDiffAE learns a disentangled latent space and is capable of generating high-quality counterfactual images.

1 Introduction

Diffusion probabilistic models (DPMs) [31, 10, 20, 32, 33] are a class of likelihood-based generative models that have achieved remarkable successes in the generation of high-resolution images with many large-scale implementations such as DALLÉ-2 [26], Stable Diffusion [27], and Imagen [28]. Thus, there has been great interest in evaluating the capabilities of diffusion models. Two of the most promising approaches are formulated as discrete-time [10] and continuous-time [33] step-wise perturbations of the data distribution. A model is then trained to estimate the reverse process which transforms noisy samples to samples from the underlying data distribution. Representation learning has been an integral component of generative models such as GANs [7] and VAEs [13] for extracting robust and interpretable features from complex data [30, 2, 25]. Recently, a thrust of research

has focused on whether DPMs can be used to extract a semantically meaningful and decodable representation that increases the quality of and control over generated images [21, 24]. However, there has been no work in modeling causal relations among the semantic latent codes to learn causal representations and enable *counterfactual generation* at inference time in DPMs. Generating high-quality counterfactual images is critical for domains such as healthcare and medicine [17, 29]. The ability to generate accurate counterfactual data from a causal graph obtained from domain knowledge can significantly cut the cost of data collection. Furthermore, reasoning about hypothetical scenarios unseen in the training distribution can be quite insightful for gauging the interactions among causal variables in complex systems. Given a causal graph of a system, we study the capability of DPMs as causal representation learners and evaluate their ability to generate counterfactuals upon interventions on causal variables.

Intuitively, we can think about the DPM as an encoder-decoder framework. The encoding maps an input image x_0 to a spatial latent variable x_T through a series of Gaussian noise perturbations. However, x_T can be interpreted as a noise representation that lacks high-level semantics [24]. Recently, Preechakul et al [24] proposed a diffusion-based autoencoder (DiffAE) to extract a high-level semantic representation alongside the stochastic low-level representation x_T for decodable representation learning. Learning such a semantic representation also enables interpolation in the latent space for controllable generation and has been shown to improve image generation quality. Mittal et al [19] built on this framework and introduced a diffusion-based representation learning (DRL) objective that instead learns time-conditioned representations throughout the diffusion process. However, both these approaches learn arbitrary representations and do not focus on disentanglement, a key property of interpretable representations. Disentangled representations enable precise control of generative factors in isolation. When considering causal systems, disentanglement is important for performing isolated interventions.

In this paper, we focus on learning disentangled causal representations, where the high-level semantic factors are causally related. To the best of our knowledge, we are the first to explore representation-based counterfactual image generation using diffusion probabilistic models. We propose **CausalDiffAE**, a learning framework for causal representation learning and controllable counterfactual generation in DPMs. Our key idea is to learn a causal representation via a learnable stochastic encoder and model the relations among latents via causal mechanisms parameterized by neural networks. We formulate a variational objective with a label alignment prior to enforce disen-

* Corresponding Author. Email: akomandu@uark.edu.

tanglement of the learned causal factors. We then utilize a conditional denoising diffusion implicit model (DDIM) [32] for decoding and modeling the stochastic variations. Intuitively, the causal representation encodes compact information that is *causally relevant* for image decoding in reverse diffusion. Furthermore, the modeling of causal relations in the latent space enables the generation of counterfactuals upon interventions on learned causal variables. We propose a DDIM variant for counterfactual generation subject to $\mathbf{do}(\cdot)$ interventions [23]. In an effort to improve the practicality and interpretability of the model, we propose an extension to CausalDiffAE that utilizes weaker supervision. In the scenario where labeled data is limited, we jointly train an unconditional and representation-conditioned diffusion model on the unlabeled and labeled partitions, respectively. This approach significantly reduces the number of labeled samples required for training and enables granular control over the strength of interventions and the quality of generated counterfactuals.

2 Related Work

Recent work in causal generative modeling has focused on either learning causal representations or controllable counterfactual generation [15]. Yang et al proposed CausalVAE [34], a causal representation learning framework that models latent causal variables by a linear SCM. Kocaoglu et al [14] proposed CausalGAN, an extension of the GAN to model causal variables for sampling from interventional distributions. Diffusion and score-based generative models [10, 33] have shown impressive results in class-conditional generation either through classifier-based [5] or classifier-free [9] paradigms. Recently, there has been an interest in exploring the capacity of diffusion models as representation learners. For instance, Mittal et al [19] and Preechakul et al [24] considered diffusion-based representation learning objectives. Mamaghan et al [11] explored representation learning from a score-based perspective given access to data in the form of counterfactual pairs. However, this work does not focus on counterfactual generation. Another related area of research is counterfactual explanations [1], which focuses on post-hoc methods to generate realistic counterfactuals, but not in the strictly causal sense. Our work focuses on diffusion-based *representation learning* and is most closely related to DiffAE [24] and DRL [19], which aim to learn semantically meaningful representations. However, the key distinction is that we learn causal representations to enable counterfactual generation. Our proposed framework extends CausalVAE to diffusion-based models and under a weaker supervision paradigm.

3 Background

3.1 Structural Causal Model

A structural causal model (SCM) is formally defined by a tuple $\mathcal{M} = \langle \mathcal{Z}, \mathcal{U}, F \rangle$, where \mathcal{Z} is the domain of the set of n endogenous causal variables $\mathbf{z} = \{z_1, \dots, z_n\}$, \mathcal{U} is the domain of the set of n exogenous noise variables $\mathbf{u} = \{u_1, \dots, u_n\}$, which is learned as an intermediate latent variable, and $F = \{f_1, \dots, f_n\}$ is a collection of n independent causal mechanisms of the form

$$z_i = f_i(u_i, z_{\mathbf{pa}_i}) \quad (1)$$

where $\forall i, f_i : \mathcal{U}_i \times \prod_{j \in \mathbf{pa}_i} \mathcal{Z}_j \rightarrow \mathcal{Z}_i$ are **causal mechanisms** that determine each causal variable as a function of the parents and noise, $z_{\mathbf{pa}_i}$ are the parents of causal variable z_i ; and a probability measure $p_{\mathcal{U}}(\mathbf{u}) = p_{\mathcal{U}_1}(u_1)p_{\mathcal{U}_2}(u_2) \dots p_{\mathcal{U}_n}(u_n)$, which admits a product distribution. For the purposes of this work, we assume a causally

sufficient setting (no hidden confounding), no SCM misspecification, and faithfulness is satisfied.

3.2 Diffusion Probabilistic Models

Diffusion Probabilistic Models (DPMs) [10, 20] have shown impressive results in image generation tasks, even beating out GANs in many cases [5]. The idea of the denoising diffusion probabilistic model (DDPM) [10] is to define a Markov chain of diffusion steps to slowly destroy the structure in a data distribution through a forward diffusion process by adding noise [10] and learn a reverse diffusion process that restores the structure of the data. Some proposed methods, such as denoising diffusion implicit model (DDIM) [32], break the Markov assumption to speed up the sampling in the diffusion process by carrying out a deterministic encoding of the noise.

Forward Diffusion. Given some input data sampled from a distribution $\mathbf{x}_0 \sim q(\mathbf{x})$, the forward diffusion process is defined by adding small amounts of Gaussian noise to the sample in T steps thereby producing noisy samples $\mathbf{x}_1, \dots, \mathbf{x}_T$. The distribution of the noisy sample at time step t is defined as a conditional distribution as follows:

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}) \quad (2)$$

where $\beta_t \in (0, 1)$ is a variance parameter that controls the step size of noise. As $t \rightarrow \infty$, the input sample \mathbf{x}_0 loses its distinguishable features. In the end, when $t = T$, \mathbf{x}_T follows an isotropic Gaussian. From Eq (2), we can then define a closed-form tractable posterior over all time steps factorized as follows:

$$q(\mathbf{x}_{1:T} | \mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1}) \quad (3)$$

Now, \mathbf{x}_t can be sampled at any arbitrary time step t using the reparameterization trick. Let $\alpha_t = \prod_{i=1}^t 1 - \beta_i$:

$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t} \mathbf{x}_0, (1 - \alpha_t) \mathbf{I}) \quad (4)$$

Reverse Diffusion. In the reverse process, to sample from $q(\mathbf{x}_{t-1} | \mathbf{x}_t)$, the goal is to recreate the true sample \mathbf{x}_0 from a Gaussian noise input $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Unlike the forward diffusion, $q(\mathbf{x}_{t-1} | \mathbf{x}_t)$ is not analytically tractable and thus requires learning a model p_θ to approximate the conditional distributions as follows:

$$p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) \quad (5)$$

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \Sigma_\theta(\mathbf{x}_t, t))$$

where μ_θ and Σ_θ are learned via neural networks. It turns out that conditioning on the input \mathbf{x}_0 yields a tractable reverse conditional probability

$$q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{t-1}; \tilde{\mu}(\mathbf{x}_t, \mathbf{x}_0), \tilde{\beta}_t \mathbf{I}) \quad (6)$$

where $\tilde{\mu}$ and $\tilde{\beta}_t$ are the true mean and variance. The learning objective is then formulated as a simplified objective of the ELBO via reparameterization to minimize the following mean squared error loss

$$\mathcal{L}_{\text{simple}} = \sum_{t=1}^T \mathbb{E}_{\mathbf{x}_0, \epsilon_t} \left[\|\epsilon_t - \epsilon_\theta(\mathbf{x}_t, t)\|_2^2 \right] \quad (7)$$

where $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is the noise that takes an analytical form via a reparameterization from \mathbf{x}_0 , as shown in [10].

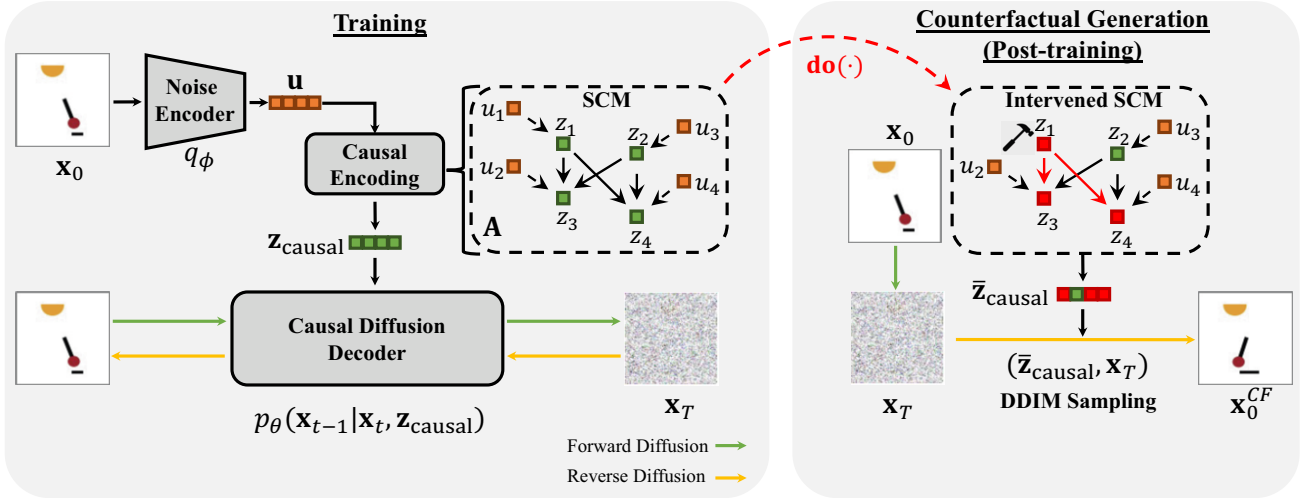


Figure 1: CausalDiffAE Framework. The left side details the training process of CausalDiffAE by encoding to causal representation $\mathbf{z}_{\text{causal}}$ and using a conditional DDIM decoder conditioned on $\mathbf{z}_{\text{causal}}$ and \mathbf{x}_T for image reconstruction. The right side shows the DDIM-based counterfactual generation procedure using a trained CausalDiffAE model.

DPMs produce latent variables $\mathbf{x}_{1:T}$ through the forward process. However, these variables are stochastic [24]. Song et al. proposed a DPM called Denoising Diffusion Implicit model (DDIM), which enables a deterministic process as follows:

$$\mathbf{x}_{t-1} = \sqrt{\alpha_{t-1}} \left(\frac{\mathbf{x}_t - \sqrt{1 - \alpha_t} \epsilon_\theta(\mathbf{x}_t, t)}{\sqrt{\alpha_t}} \right) + \sqrt{1 - \alpha_{t-1}} \epsilon_\theta(\mathbf{x}_t, t) \quad (8)$$

with the following deterministic decoding process

$$q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) = \mathcal{N} \left(\sqrt{\alpha_{t-1}} \mathbf{x}_0 + \sqrt{1 - \alpha_{t-1}} \frac{\mathbf{x}_t - \sqrt{\alpha_t} \mathbf{x}_0}{\sqrt{1 - \alpha_t}}, \mathbf{0} \right) \quad (9)$$

which keeps the DDPM marginal distribution $q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\sqrt{\alpha_t} \mathbf{x}_0, (1 - \alpha_t) \mathbf{I})$. It turns out that this formulation shares the same objective and solution of DDPM and only differs in the sampling procedure. Thus, we can *deterministically* obtain the noise map \mathbf{x}_T corresponding to a given image \mathbf{x}_0 .

4 Causal Diffusion Autoencoders

Existing diffusion-based controllable generation methods neglect the scenario where generative factors are causally related and do not support counterfactual generation. To tackle this issue, we propose CausalDiffAE, a diffusion-based causal representation learning framework to enable counterfactual generation. Firstly, we define a latent SCM to describe the semantic causal representation as a function of learned noise encodings. In the case of diffusion autoencoders [24], the semantic latent representation \mathbf{z}_{sem} captures high-level semantic information, and \mathbf{x}_T captures low-level stochastic information. In our formulation, we learn a causal representation $\mathbf{z}_{\text{causal}}$ which captures *causally relevant* information. Together, the two latent variables ($\mathbf{z}_{\text{causal}}, \mathbf{x}_T$) capture all the detailed causal semantics and stochasticity in the image. Secondly, given a trained CausalDiffAE model, we propose a counterfactual generation algorithm that utilizes $\text{do}(\cdot)$ interventions and the DDIM sampling algorithm. The overall framework of CausalDiffAE is shown in Figure 1.

4.1 Causal Encoding

Let $\mathbf{x}_0 \in \mathbb{R}^d$ be the observed input image. We carry out the forward diffusion process until we have a set of T perturbed samples $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$, each at a different noise scale. Suppose there are n abstract causal variables that describe the high-level semantics of the observed image. To learn a meaningful representation, we propose to encode the input image \mathbf{x}_0 to a low-dimensional noise encoding $\mathbf{u} \in \mathbb{R}^n$. We then map the noise encoding to latent causal factors $\mathbf{z}_{\text{causal}} \in \mathbb{R}^n$ corresponding to the abstract causal variables. In this formulation, each noise term u_i is the exogenous noise term for causal variable z_i in the SCM. Let \mathbf{A} be the adjacency matrix encoding the causal graph among the underlying factors where A_{ji} implies z_j is a cause of z_i . Then, we parameterize the mechanisms among causal variables as follows

$$z_i = f_i(z_{\text{pa}_i}, u_i) \quad (10)$$

where f_i is the causal mechanism generating causal variable z_i as a function of its parents and exogenous noise term and z_{pa_i} denotes the causal parents of factor z_i . In practice, we can implement f_i as a post-nonlinear additive noise model such that

$$\mathbf{z} = (\mathbf{I} - \mathbf{A}^T)^{-1} \mathbf{u} \quad (11)$$

$$z_i = f_i(\mathbf{A}_i \odot \mathbf{z}; \nu_i) + u_i$$

where ν_i are the parameters of the neural network parameterizing each mechanism, \odot is the elementwise product, and $\mathbf{z}_{\text{causal}} = \{z_1, \dots, z_n\}$. This module captures the causal relations between latent variables using neural structural causal models. For the purposes of this work, we assume that the causal graph is known since we focus on counterfactual generation. However, a more end-to-end framework may include a causal discovery component. See Appendix C in [16] for a more detailed discussion.

4.2 Generative Model

Let \mathbf{x}_0 denote the high-dimensional input image and $\mathbf{y} \in \mathbb{R}^n$ denote an auxiliary weak supervision signal. Then, the CausalDiffAE generative process can be factorized as follows:

Algorithm 1 CausalDiffAE Training**Input:** (image, label) pairs $(\mathbf{x}_0, \mathbf{y})$ **Output:** learned parameters $\{\theta, \phi\}$

-
- 1:
- repeat**
-
- 2:
- $\mathbf{x}_0 \sim q(\mathbf{x}_0)$
-
- 3:
- $\mathbf{u} \sim q_\phi(\mathbf{u}|\mathbf{x}_0)$
- ▷ Noise encoding
-
- 4:
- $\mathbf{z}_{\text{causal}} = \{f_i(u_i, z_{\text{pa}_i}; \nu_i)\}_{i=1}^n$
- ▷ Causal encoding
-
- 5:
- $t \sim \mathcal{U}(\{1, \dots, T\})$
- ▷ Sample timestep
-
- 6:
- $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
-
- 7:
- $\mathbf{x}_t = \sqrt{\alpha_t}\mathbf{x}_0 + \sqrt{1 - \alpha_t}\epsilon_t$
- ▷ Corrupt data to sampled time
-
- 8: Take gradient step on
- $\nabla_{\theta, \phi} \mathcal{L}_{\text{CausalDiffAE}}$
-
- 9:
- until**
- convergence
-

$$p(\mathbf{x}_{0:T}, \mathbf{u}, \mathbf{z}_{\text{causal}}|\mathbf{y}) = p_\theta(\mathbf{x}_{0:T}|\mathbf{u}, \mathbf{z}_{\text{causal}}, \mathbf{y})p(\mathbf{u}, \mathbf{z}_{\text{causal}}|\mathbf{y}) \quad (12)$$

where θ are the parameters of the reverse process of the causal diffusion decoder (will discuss in Section 4.3), $p(\mathbf{u}, \mathbf{z}_{\text{causal}}|\mathbf{y}) = p(\mathbf{u})p(\mathbf{z}_{\text{causal}}|\mathbf{y})$, $p(\mathbf{u}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$, and $p(\mathbf{z}_{\text{causal}}|\mathbf{y})$ is the alignment prior defined in Eq. (19). The joint posterior distribution $p(\mathbf{x}_{1:T}, \mathbf{u}, \mathbf{z}_{\text{causal}}|\mathbf{x}_0, \mathbf{y})$ is intractable, so we approximate it using a variational distribution $q(\mathbf{x}_{1:T}, \mathbf{u}, \mathbf{z}_{\text{causal}}|\mathbf{x}_0, \mathbf{y})$ which can be factorized into the following conditional distributions

$$q(\mathbf{x}_{1:T}, \mathbf{u}, \mathbf{z}_{\text{causal}}|\mathbf{x}_0, \mathbf{y}) = q_\phi(\mathbf{z}_{\text{causal}}, \mathbf{u}|\mathbf{x}_0, \mathbf{y}) \quad (13)$$

$$q(\mathbf{x}_{1:T}|\mathbf{u}, \mathbf{z}_{\text{causal}}, \mathbf{x}_0)$$

where ϕ are the parameters of the variational encoder network parameterizing the joint distribution over the noise \mathbf{u} and causal factors $\mathbf{z}_{\text{causal}}$. We can remove the dependence on \mathbf{y} for the second conditional term in the decomposition of Eq. (13) since $\mathbf{x}_{1:T}$ is independent of the auxiliary label \mathbf{y} . We note that $q_\phi(\mathbf{z}_{\text{causal}}, \mathbf{u}|\mathbf{x}_0, \mathbf{y})$ can be factorized as $q_\phi(\mathbf{z}_{\text{causal}}|\mathbf{x}_0, \mathbf{y})q_\phi(\mathbf{u}|\mathbf{x}_0)$ since \mathbf{u} and $\mathbf{z}_{\text{causal}}$ have a one-to-one correspondence.

4.3 Causal Diffusion Decoder

We use a conditional DDIM decoder that takes as input the pair of latent variables $(\mathbf{z}_{\text{causal}}, \mathbf{x}_T)$ to generate the output image. We approximate the inference distribution $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ by parameterizing the probabilistic decoder via a conditional DDIM $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{z}_{\text{causal}})$. With DDIM, the forward process becomes completely deterministic except for $t = 1$. Similar to [24], we define the joint distribution of the reverse generative process as follows:

$$p_\theta(\mathbf{x}_{0:T}|\mathbf{z}_{\text{causal}}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{z}_{\text{causal}}) \quad (14)$$

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{z}_{\text{causal}}) = \begin{cases} \mathcal{N}(\mathbf{f}_\theta(\mathbf{x}_1, 1, \mathbf{z}_{\text{causal}}), \mathbf{0}) & \text{if } t = 1 \\ q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{f}_\theta(\mathbf{x}_t, t, \mathbf{z}_{\text{causal}})) & \text{otherwise} \end{cases} \quad (15)$$

where \mathbf{f}_θ is parameterized by a noise prediction network ϵ_θ (i.e., UNet [5]) as follows:

$$\mathbf{f}_\theta(\mathbf{x}_t, t, \mathbf{z}_{\text{causal}}) = \frac{1}{\sqrt{\alpha_t}}(\mathbf{x}_t - \sqrt{1 - \alpha_t}\epsilon_\theta(\mathbf{x}_t, t, \mathbf{z}_{\text{causal}})) \quad (16)$$

Note that in Eq. (14), \mathbf{u} is omitted since $\mathbf{z}_{\text{causal}}$ already captures all the information about the noise. By leveraging the reparameterization

trick, we can optimize the following mean squared error between noise terms

$$\mathcal{L}_{\text{simple}} = \sum_{t=1}^T \mathbb{E}_{\mathbf{x}_0, \epsilon_t} \left[\|\epsilon_\theta(\mathbf{x}_t, t, \mathbf{z}_{\text{causal}}) - \epsilon_t\|_2^2 \right] \quad (17)$$

where $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\mathbf{x}_t = \sqrt{\alpha_t}\mathbf{x}_0 + \sqrt{1 - \alpha_t}\epsilon_t$.

4.4 Learning Objective

To ensure the causal representation is disentangled, we incorporate label information $\mathbf{y} \in \mathbb{R}^n$ as a prior in the variational objective to aid in learning semantic factors and for identifiability guarantees [12]. We define the following joint loss objective:

$$\mathcal{L}_{\text{CausalDiffAE}} = \mathcal{L}_{\text{simple}} + \gamma \left\{ \mathcal{D}_{KL}(q_\phi(\mathbf{z}_{\text{causal}}|\mathbf{x}_0, \mathbf{y})\|p(\mathbf{z}_{\text{causal}}|\mathbf{y})) + \mathcal{D}_{KL}(q_\phi(\mathbf{u}|\mathbf{x}_0)\|\mathcal{N}(\mathbf{0}, \mathbf{I})) \right\} \quad (18)$$

where γ is a regularization hyperparameter similar to the bottleneck parameter in β -VAEs [8], and the alignment prior over latent variables is defined as the following exponential family distribution

$$p(\mathbf{z}_{\text{causal}}|\mathbf{y}) = \prod_{i=1}^n p(z_i|y_i) = \prod_{i=1}^n \mathcal{N}(z_i; \mu_\nu(y_i), \sigma_\nu^2(y_i)\mathbf{I}) \quad (19)$$

where μ_ν and σ_ν^2 are functions that estimate the mean and variance of the Gaussian, respectively. Intuitively, this prior ensures that the learned factors are one-to-one mapped to an indicator of the underlying ground truth factors. DiffAE requires training a latent DDIM in the latent space of the pre-trained autoencoder to enable sampling of latent semantic representation. However, CausalDiffAE is formulated as a variational objective with a stochastic encoder. Thus, we can sample the representation from the defined prior directly without having to train a separate diffusion model in the latent space. The training procedure for CausalDiffAE is outlined in Algorithm 1. See Appendix A in [16] for a derivation of the ELBO. For a discussion on the connection to score-based generative models [33], see Appendix B in [16].

4.5 Counterfactual Generation

A fundamental property of causal models is the ability to perform interventions and observe changes to a system. In generative models, this enables the sampling of counterfactual data. Given a pre-trained CausalDiffAE, we can controllably manipulate any factor of variation, propagate the causal effects to descendants, and perform reverse diffusion to sample from the counterfactual distribution. Algorithm 2 shows the process of generating counterfactuals from a trained CausalDiffAE, where \mathbf{x}_0 refers to the factual observation and \mathbf{x}_0^{CF} refers to the generated counterfactual sample. To generate counterfactual instances, we first encode the high dimensional observation \mathbf{x}_0 to a noise encoding \mathbf{u} (abduction) and transform it to causal latent variables $\mathbf{z}_{\text{causal}}$. Then, we intervene on a desired variable and propagate the causal effects via neural mechanisms to yield the intervened representation $\mathbf{z}_{\text{causal}}$. We utilize the DDIM sampling algorithm to ensure the stochastic noise \mathbf{x}_T is a deterministic encoding to enable semantic manipulations. Finally, we decode using DDIM conditioned

Algorithm 2 CausalDiffAE Counterfactual Generation

Input: Factual sample \mathbf{x}_0 , intervention target set \mathcal{I} with intervention values c , noise predictor ϵ_θ , encoder ϕ
Output: Counterfactual sample \mathbf{x}_0^{CF}

```

1:  $\mathbf{u} \sim q_\phi(\mathbf{u}|\mathbf{x}_0)$  ▷ Noise encoding
2: for  $i = 1$  to  $n$  do ▷ in topological order
3:   if  $i \in \mathcal{I}$  then
4:      $z_i = c_i$ 
5:   else
6:      $z_i = f_i(u_i, z_{\text{pa}_i})$ 
7:   end if
8: end for
9:  $\bar{\mathbf{z}}_{\text{causal}} = \{z_1, \dots, z_n\}$  ▷ Intervened representation
10:  $\mathbf{x}_T \sim \mathcal{N}(\sqrt{\alpha_T}\mathbf{x}_0, (1 - \alpha_T)\mathbf{I})$ 
11:  $\mathbf{x}_T^{CF} = \mathbf{x}_T$ 
12: for  $t = T, \dots, 1$  do ▷ DDIM sampling
13:    $\mathbf{x}_{t-1}^{CF} = \sqrt{\alpha_{t-1}} \left( \frac{\mathbf{x}_t^{CF} - \sqrt{1 - \alpha_t} \epsilon_\theta(\mathbf{x}_t^{CF}, t, \bar{\mathbf{z}}_{\text{causal}})}{\sqrt{\alpha_t}} \right)$ 
14:    $\quad + \sqrt{1 - \alpha_{t-1}} \epsilon_\theta(\mathbf{x}_t^{CF}, t, \bar{\mathbf{z}}_{\text{causal}})$ 
15: end for
16: return  $\mathbf{x}_0^{CF}$ 

```

on $(\bar{\mathbf{z}}_{\text{causal}}, \mathbf{x}_T)$ to obtain a counterfactual \mathbf{x}_0^{CF} . In lines 12-13, we use the DDIM non-Markovian deterministic generative process to generate counterfactual instances as follows:

$$\mathbf{x}_{t-1}^{CF} = \sqrt{\alpha_{t-1}} \left(\frac{\mathbf{x}_t^{CF} - \sqrt{1 - \alpha_t} \epsilon_\theta(\mathbf{x}_t^{CF}, t, \bar{\mathbf{z}}_{\text{causal}})}{\sqrt{\alpha_t}} \right) + \sqrt{1 - \alpha_{t-1}} \epsilon_\theta(\mathbf{x}_t^{CF}, t, \bar{\mathbf{z}}_{\text{causal}}) \quad (20)$$

Conditioning vs. Intervening. When we study causal generative models, we utilize the intervention operation, which is a fundamentally different operation than conditioning. When we condition, we narrow our scope to a specific subgroup of the data based on the conditioning variable. Interventions are population-level operations that fix a variable’s value (rendering it independent of its parents) to determine causal effects downstream. We emphasize that, under this intervention operation, causal models are robust to distribution shifts and can generate data outside the support of the training distribution.

4.6 Weak Supervision

To reduce the reliance on labeled data, inspired by classifier-free [9] guidance, we train a CausalDiffAE with a weak supervision guidance paradigm on the representation level [5].

Training. In the limited labeled-data regime, we train two models: an unconditional denoising diffusion model $p_\theta(\mathbf{x})$ parameterized by the score estimator $\epsilon_\theta(\mathbf{x}_t, t)$ and a representation-conditioned model $p_\theta(\mathbf{x}|\mathbf{z}_{\text{causal}})$ parameterized through $\epsilon_\theta(\mathbf{x}_t, t, \mathbf{z}_{\text{causal}})$. We use a single neural network to parameterize both models, where for the unconditional model we use only the unlabeled data for predicting the score (i.e., $\epsilon_\theta(\mathbf{x}_t, t)$).

Generation. The counterfactual generation procedure in lines 12-13 of Algorithm 2 can be modified to generate counterfactuals with a guidance strength ω , which can be interpreted as controlling the strength of the intervention on the causal variable to generate the counterfactual in our case. The overall modified score estimation during generation can be performed using the following linear combination of conditional and unconditional score estimates

$$\bar{\epsilon}_\theta(\mathbf{x}_t, t, \bar{\mathbf{z}}_{\text{causal}}) = \underbrace{\omega \epsilon_\theta(\mathbf{x}_t, t, \bar{\mathbf{z}}_{\text{causal}})}_{\text{causal conditional model}} + \underbrace{(1 - \omega) \epsilon_\theta(\mathbf{x}_t, t)}_{\text{unconditional model}} \quad (21)$$

where $\bar{\mathbf{z}}_{\text{causal}}$ is the set of latent causal factors after an intervention. The original utility of the classifier-free paradigm was to decrease the generation of diverse data in favor of higher-quality image samples without needing classifier gradients. So, ω controls the trade-off between higher quality and diverse samples. In our case, we care about generating high-quality counterfactual data. Intuitively, a higher ω implies a stronger effect of the intervention on the generated counterfactual since the conditional model $\epsilon_\theta(\mathbf{x}_t, t, \mathbf{z}_{\text{causal}})$ is sensitive to interventions. So, as ω decreases, the unconditional model dilutes the effect of the intervention-sensitive model. In this sense, the sampling mechanism can be used to evaluate the causal strength of interventions. We find that the weak supervision paradigm enables (1) more efficient training with a weaker supervision signal, and (2) fine-grained control over generated counterfactuals.

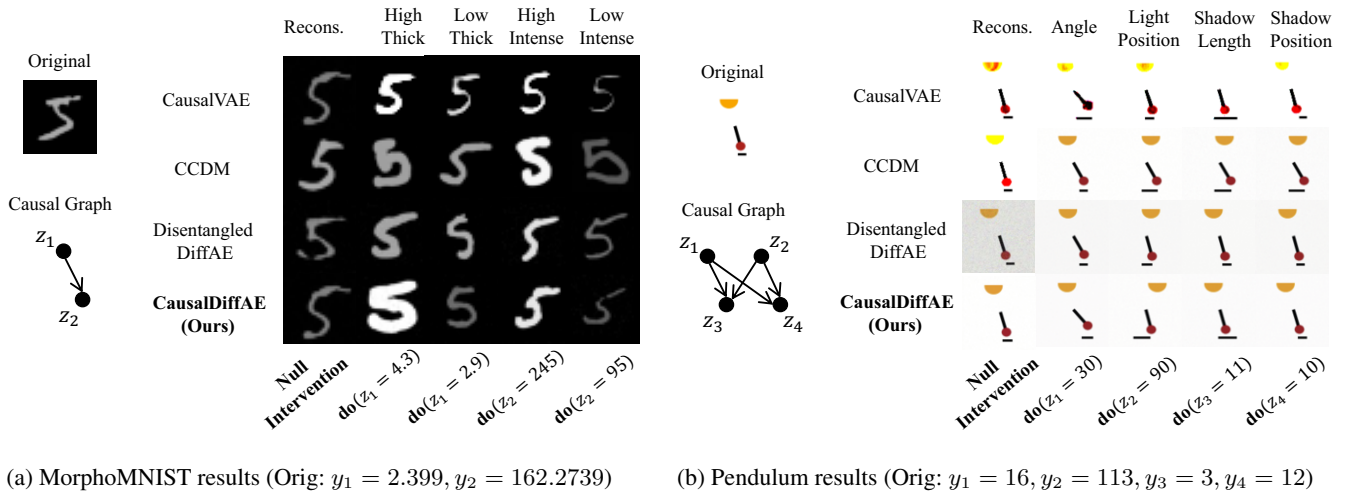
5 Experiments

5.1 Empirical Setting

Datasets. We experiment on three datasets. We use the MorphoMNIST dataset [4] created by imposing a 2-variable SCM to generate morphological transformations on the original MNIST dataset, where thickness is the cause of the intensity of the digit [22], as shown in Figure 2a. The Pendulum dataset [34] consists of images of a causal system consisting of a light source, pendulum, and shadow. The light source and the pendulum angle determine the length and position of the shadow, as shown in the causal graph in Figure 2b. We also use CausalCircuit [3], a complex 3D robotics dataset where a robot arm moves around to turn on red, green, or blue lights. The causal graph of this system is shown in Figure 3.

Baselines. CausalVAE [34] is a VAE-based causal representation learning framework that models causal variables using a linear SCM and enables counterfactual generation during inference time through interventions on causal variables. Class-conditional diffusion model (CCDM) [20] is a conditional diffusion model that utilizes class labels as the conditioning signal in reverse diffusion. Thus, this model is capable of generating new samples determined by a discrete or continuous set of labels \mathbf{y} . DiffAE [24] is a diffusion model that aims to learn manipulable and semantically meaningful latent codes. However, this approach learns an arbitrary representation in an unsupervised fashion and does not disentangle the latent space. Manipulations are performed using a post-hoc classifier for linear interpolation. Thus, the learned representation would not be ideal to perform causal interventions. For a fair comparison in counterfactual generation, we modify the objective to disentangle the latent space by incorporating label information in a prior to regularize the posterior. We call this extension **DisDiffAE**. We use DisDiffAE as a baseline to evaluate counterfactual generation and the DiffAE to evaluate disentanglement.

Metrics. We primarily use two quantitative metrics to evaluate the performance of our approach. To evaluate the disentanglement of the learned representations, we use the *DCI disentanglement* metric [6]. A high DCI score also suggests the effectiveness of controllable generation. In the context of a causal representation, this means that we can intervene on latent codes in an isolated fashion without any entanglements (i.e., two factors are encoded in the same latent code). To quantitatively evaluate generated counterfactuals, we adopt the *Effectiveness* metric from Melistas et al [18], which evaluates how

(a) MorphoMNIST results (Orig: $y_1 = 2.399, y_2 = 162.2739$)(b) Pendulum results (Orig: $y_1 = 16, y_2 = 113, y_3 = 3, y_4 = 12$)**Figure 2:** Counterfactual trajectories generated by CausalDiffAE and baseline models for (a) MorphoMNIST and (b) Pendulum datasets. We observe that CausalDiffAE generates much more accurate counterfactuals upon interventions on causal factors compared to baselines.**Table 1:** Disentanglement (DCI)

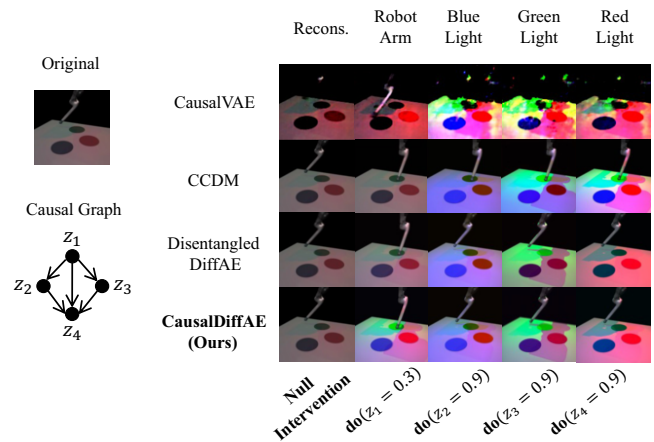
Dataset	Model	DCI \uparrow
MorphoMNIST	CausalVAE	0.784 ± 0.01
	DiffAE	0.358 ± 0.01
	CausalDiffAE	0.993 ± 0.01
Pendulum	CausalVAE	0.885 ± 0.01
	DiffAE	0.353 ± 0.01
	CausalDiffAE	0.999 ± 0.01
CausalCircuit	CausalVAE	0.8860 ± 0.01
	DiffAE	0.353 ± 0.01
	CausalDiffAE	0.999 ± 0.01

successful the performed intervention was at generating the counterfactual. We train anti-causal predictors via convolutional regressors on the training dataset for each continuous causal variable z_i . Then, we report the mean absolute error (MAE) loss between the predicted values from the generated counterfactual and the true values of the counterfactual. This metric captures how controllable the factors are and the accuracy of the generated counterfactuals.

For details about the datasets, implementation, metrics, and computational requirements, see Appendix D in [16]. Our code is available at <https://github.com/Akomand/CausalDiffAE>.

5.2 Disentanglement of Latent Space

We compare the disentanglement of CausalDiffAE with other baseline models, as shown in Table 1. We observe that diffusion-based representation learning objectives coupled with a suitable prior can better disentangle latent variables compared to VAE-based models. We do not include CCDM as a baseline here since it does not produce a representation to be evaluated. Compared to CausalVAE, the diffusion-based decoder in CausalDiffAE disentangles the semantic factors of variation to a much greater degree. Thus, we can perform interventions on causal variables in isolation and observe their downstream effects. We also note that DiffAE [24] does not learn a disentangled latent space since the semantic representation learned is arbitrary. To perform controllable manipulations with DiffAE, a post-hoc classifier must be trained to guide the sampling process. CausalDiffAE offers more precise control over learned factors through the disentanglement objective without the need to train additional classifiers.

**Figure 3:** CausalCircuit results (Orig: $y_1 = 0.02, y_2 = 0.03, y_3 = 0.04, y_4 = 0.14$)

5.3 Controllable Counterfactual Generation

Qualitative Evaluation. We show that CausalDiffAE produces much more realistic counterfactual samples compared to other acausal baselines and its VAE counterpart, CausalVAE. We attribute this to the diffusion process, which is better capable of capturing causally relevant information along with low-level stochastic variation.

Figure 2a shows the counterfactual generation results for the MorphoMNIST dataset. CausalVAE can generate counterfactual images after intervening on either thickness or intensity, but the accuracy and quality of the generated counterfactuals is far lower than CausalDiffAE. For instance, lower thickness does not lower the intensity and lower intensity intervention seems to change the thickness of the digit. CCDM fails to produce samples consistent with the underlying causal model. For example, intervening on the intensity produces a sample that increases in thickness. From a conditioning perspective, high-intensity digits tend to be thicker in the training distribution. For DisDiffAE, increasing the thickness does not influence the intensity.

Figure 2b shows counterfactual generation results for the Pendulum causal system. Upon interventions, images generated by CCDM are not consistent with the causal model. For example, intervening on the light position changes both the light position and the pendulum angle. DisDiffAE produces images where we can control one factor at

Table 2: Effectiveness on MorphoMNIST test set (MAE)

Factor	Model	Intervention	
		$do(t)$	$do(i)$
Thickness (t)	CausalVAE	3.763 ± 0.01	4.645 ± 0.01
	DisDiffAE	0.377 ± 0.02	0.326 ± 0.02
	CausalDiffAE	0.392 ± 0.02	0.309 ± 0.02
Intensity (i)	CausalVAE	13.233 ± 0.01	15.087 ± 0.01
	DisDiffAE	0.794 ± 0.02	0.262 ± 0.02
	CausalDiffAE	0.503 ± 0.01	0.256 ± 0.01

Table 3: Effectiveness on Pendulum test set (MAE)


Factor	Model	Intervention			
		$do(a)$	$do(lp)$	$do(sl)$	$do(sp)$
Angle (a)	CausalVAE	24.860	23.030	20.470	11.580
	DisDiffAE	0.668	0.648	0.647	0.647
	CausalDiffAE	0.297	0.132	0.031	0.034
LightPos (lp)	CausalVAE	34.200	26.010	35.490	47.060
	DisDiffAE	0.656	0.654	0.630	0.651
	CausalDiffAE	0.045	0.434	0.035	0.064
ShadowLen (sl)	CausalVAE	1.946	1.43	2.02	1.72
	DisDiffAE	0.550	0.527	0.560	0.516
	CausalDiffAE	0.136	0.322	0.492	0.082
ShadowPos (sp)	CausalVAE	52.52	72.50	57.03	32.78
	DisDiffAE	0.474	0.475	0.479	0.534
	CausalDiffAE	0.146	0.303	0.064	0.471

* Standard error is roughly in the range ± 0.01 to ± 0.02 for all averages.

a time, but does not reflect causal effects. For example, changing the angle of the pendulum does not accurately change the shadow length and position. CausalDiffAE generates higher-quality counterfactuals that are consistent with the causal model. Specifically, intervention on the pendulum angle or light position changes the shadow length and position accurately. On the other hand, interventions on children variables leave the parents unchanged.

Figure 3 shows counterfactuals from the CausalCircuit dataset. CausalVAE generates inaccurate counterfactuals for many scenarios (e.g., intervention on the blue light intensity changes the intensity of all other lights). For CCDM, moving the robot arm over the green light fails to turn the light on. Furthermore, manipulating the light intensity of other lights affects the position of the robot arm. DisDiffAE enables control over the generative factors, but does not consider causal effects. For example, moving the robot arm over the green light does not turn it on. Counterfactuals generated from CausalDiffAE are consistent with the causal system. For example, moving the robot arm over the green button turns the light on and as a result also turns on the red light, which is a downstream child variable. Intervening on the blue or green light slightly increases the intensity of the red light. Intervening on the red light leaves all parent variables unchanged. For additional counterfactual generation results, see Appendix D.5 in [16].

Quantitative Evaluation. We quantitatively show using the effectiveness metric that CausalDiffAE generates counterfactuals that are both accurate and realistic. We perform random interventions from a uniform distribution over the test dataset for each causal variable. We find that CausalDiffAE almost always outperforms other baselines in the effectiveness metric, as shown in Tables 2 and 3, for all causal factors. Specifically, for the MorphoMNIST dataset, we observe that interventions on thickness produce counterfactuals that accurately reflect both the thickness and intensity values. In the scenario where we intervene on thickness, the intensity MAE is lower for CausalDiffAE than other baselines, which indicates that the generated counterfactual has an accurate intensity value consistent with the causal effect of thickness on intensity. When we intervene on intensity, the thickness

**Figure 4:** MorphoMNIST Weak Supervision

MAE is lower for CausalDiffAE than baselines, which suggests that the generated counterfactual retains its original thickness value upon intervention on intensity. For the Pendulum dataset, we see a similar phenomenon, where interventions on causal factors along with their downstream effects are accurately captured in the generated counterfactuals. We do not evaluate effectiveness for the CausalCircuit dataset since we do not have access to the generative process used to obtain the factors. We compute the average effectiveness value over 5 runs with different random seeds. Our results strongly imply that the generated counterfactuals closely match the true counterfactuals.

5.4 Case Study: Weak Supervision Results

Unlike VAE-based approaches, the weak supervision paradigm of diffusion models reduces the full-label supervision. We study the weak supervision scenario with the MorphoMNIST dataset. We jointly train a representation-conditioned and unconditional model, where the conditioned split is far less than the unconditioned split. We have two main motivations for doing this: (1) it greatly reduces the need for fully labeled datasets, and (2) it enables granular control over generated counterfactuals. We denote the proportion of unlabeled data by $p_{\text{unlabeled}}$. The CausalDiffAE model trained with $p_{\text{unlabeled}} = 0.8$ on the MorphoMNIST dataset yields a DCI score of 0.9964, which suggests that even under strictly limited label supervision, CausalDiffAE learns disentangled representations. We also empirically show that changing the ω parameter controls the strength of the intervention on the generated counterfactual. Figure 4 shows MNIST digits generated using the joint estimated score from the reduced supervision version of CausalDiffAE. We observe that interventions have virtually no effect when sampling using the joint score with $\omega = 0.2$. For $\omega = 0.5$, we see a stronger effect of the intervention on the thickness and intensity of the digit. Finally, for the fully-supervised score $\omega = 1.0$, the intervention acts the strongest. Thus, varying ω in the range $(0, 1)$ can be interpreted as generating a range of different counterfactuals.

6 Conclusion

In this work, we propose CausalDiffAE, a diffusion-based framework for causal representation learning and counterfactual generation. We propose a causal encoding mechanism that maps images to causally related factors. We learn the causal mechanisms among factors via neural networks. We formulate a variational diffusion-based objective to enforce the disentanglement of the latent space to enable latent space manipulations. We propose a DDIM-based counterfactual generation algorithm subject to interventions. For limited supervision scenarios, we propose a weak supervision extension of our model, which jointly learns an unconditional and conditional model. This objective also enables granular control over generated counterfactuals. We empirically show the capability of our model using both qualitative and quantitative metrics. Future work includes exploring counterfactual generation in text-to-image diffusion models.

Acknowledgements

This work is supported in part by National Science Foundation under awards 1910284, 1946391 and 2147375, the National Institute of General Medical Sciences of National Institutes of Health under award P20GM139768, and the Arkansas Integrative Metabolic Research Center at University of Arkansas.

References

- [1] M. Augustin, V. Boreiko, F. Croce, and M. Hein. Diffusion visual counterfactual explanations. In *Advances in Neural Information Processing Systems*, 2022.
- [2] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013.
- [3] J. Brehmer, P. D. Haan, P. Lippe, and T. Cohen. Weakly supervised causal representation learning. In *Advances in Neural Information Processing Systems*, 2022.
- [4] D. C. Castro, J. Tan, B. Kainz, E. Konukoglu, and B. Glocker. MorphoMNST: Quantitative assessment and diagnostics for representation learning. *Journal of Machine Learning Research*, 20(178), 2019.
- [5] P. Dhariwal and A. Q. Nichol. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems*, 2021.
- [6] C. Eastwood, A. L. Nicoliciou, J. V. Kügelgen, A. Kekić, F. Träuble, A. Dittadi, and B. Schölkopf. DCI-ES: An extended disentanglement framework with connections to identifiability. In *The Eleventh International Conference on Learning Representations*, 2023.
- [7] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014.
- [8] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. beta-VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2017.
- [9] J. Ho and T. Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.
- [10] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020.
- [11] A. M. Karimi Mamaghan, A. Dittadi, S. Bauer, K. H. Johansson, and F. Quinzan. Diffusion-based causal representation learning. *Entropy*, 26(7), 2024.
- [12] I. Khemakhem, D. Kingma, R. Monti, and A. Hyvarinen. Variational autoencoders and nonlinear ica: A unifying framework. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, 2020.
- [13] D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.
- [14] M. Kocaoglu, C. Snyder, A. G. Dimakis, and S. Vishwanath. CausalGAN: Learning causal implicit generative models with adversarial training. In *International Conference on Learning Representations*, 2018.
- [15] A. Komanduri, X. Wu, Y. Wu, and F. Chen. From identifiable causal representations to controllable counterfactual generation: A survey on causal generative modeling. *Transactions on Machine Learning Research*, 2024.
- [16] A. Komanduri, C. Zhao, F. Chen, and X. Wu. Causal diffusion autoencoders: Toward counterfactual generation via diffusion probabilistic models. *arXiv preprint arXiv:2404.17735*, 2024. Full version of this paper.
- [17] X. Liu, P. Sanchez, S. Thermos, A. Q. O’Neil, and S. A. Tsafaris. Learning disentangled representations in the imaging domain. *Medical Image Analysis*, 80:102516, 2022.
- [18] T. Melistas, N. Spyrou, N. Gkouti, P. Sanchez, A. Vlontzos, G. Papanastasiou, and S. A. Tsafaris. Benchmarking counterfactual image generation. *arXiv preprint arXiv:2403.20287*, 2024.
- [19] S. Mittal, K. Abstreiter, S. Bauer, B. Schölkopf, and A. Mehrjou. Diffusion based representation learning. In *Proceedings of the 40th International Conference on Machine Learning*, 2023.
- [20] A. Q. Nichol and P. Dhariwal. Improved denoising diffusion probabilistic models. In *Proceedings of the 38th International Conference on Machine Learning*, 2021.
- [21] K. Pandey, A. Mukherjee, P. Rai, and A. Kumar. DiffuseVAE: Efficient, controllable and high-fidelity generation from low-dimensional latents. *Transactions on Machine Learning Research*, 2022.
- [22] N. Pawlowski, D. Coelho de Castro, and B. Glocker. Deep structural causal models for tractable counterfactual inference. In *Advances in Neural Information Processing Systems*, 2020.
- [23] J. Pearl. *Causality*. Cambridge University Press, Cambridge, UK, 2 edition, 2009. ISBN 978-0-521-89560-6.
- [24] K. Preechakul, N. Chatthee, S. Wizadwongsa, and S. Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [25] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, 2021.
- [26] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [27] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [28] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, B. K. Ayan, S. S. Mahdavi, R. G. Lopes, T. Salimans, J. Ho, D. J. Fleet, and M. Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems*, 2022.
- [29] P. Sanchez, J. P. Voisey, T. Xia, H. I. Watson, A. Q. O’Neil, and S. A. Tsafaris. Causal machine learning for healthcare and precision medicine. *Royal Society Open Science*, 2022.
- [30] B. Scholkopf, F. Locatello, S. Bauer, N. R. Ke, N. Kalchbrenner, A. Goyal, and Y. Bengio. Toward Causal Representation Learning. *Proceedings of the IEEE*, 109:612–634, May 2021. ISSN 0018-9219, 1558-2256.
- [31] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning*, 2015.
- [32] J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021.
- [33] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- [34] M. Yang, F. Liu, Z. Chen, X. Shen, J. Hao, and J. Wang. Causalvae: Disentangled representation learning via neural structural causal models. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.