









Enhanced Controllability of Diffusion Models via Feature Disentanglement and Realism-Enhanced Sampling Methods

Wonwoong Cho^{1*}, Hareesh Ravi², Midhun Harikumar², Vinh Khuc²,
Krishna Kumar Singh³, Jingwan Lu³, David Inouye¹, and Ajinkya Kale²

¹ Purdue University, West Lafayette IN 47907, USA

² Adobe Applied Research, San Jose CA 95110, USA

³ Adobe Research, San Jose CA 95110, USA

Abstract. As Diffusion Models have shown promising performance, a lot of efforts have been made to improve the controllability of Diffusion Models. However, how to train Diffusion Models to have the disentangled latent spaces and how to naturally incorporate the disentangled conditions during the sampling process have been underexplored. In this paper, we present a training framework for feature disentanglement of Diffusion Models (FDiff). We further propose two sampling methods that can boost the realism of our Diffusion Models and also enhance the controllability. Concisely, we train Diffusion Models conditioned on two latent features, a spatial content mask, and a flattened style embedding. We rely on the inductive bias of the denoising process of Diffusion Models to encode pose/layout information in the content feature and semantic/style information in the style feature. Regarding the sampling methods, we first generalize Composable Diffusion Models (GCDM) by breaking the conditional independence assumption to allow for some dependence between conditional inputs, which is shown to be effective in realistic generation in our experiments. Second, we propose timestep-dependent weight scheduling for content and style features to further improve the performance. We also observe better controllability of our proposed methods compared to existing methods in image manipulation and image translation.

Keywords: Diffusion Models · Feature Disentanglement · Classifier-free Guidance

1 Introduction

Diffusion Models [14, 36] have gained a lot of attention due to their impressive performance in image generation [10, 31, 32] and likelihood estimation [27]. Their applications have been widely and deeply explored as well, such as conditional image generation [26, 44], personalization [11, 34], and image translation [21].

* Performed this work during internship at Adobe Applied Research.

Even though numerous research works have been studied in Diffusion Models, it has remained underexplored how to train Diffusion Models to have disentangled latent spaces, and how to sample from the Diffusion Models while incorporating the decomposed conditions naturally in a controllable way.

The idea of training generative models with multiple external disentangled latent spaces has been widely explored in GANs [9, 17, 20, 22, 28]. A common strategy across such methods is to learn a structure/content feature z_c to capture the underlying structure (e.g., facial shape and pose in face images) and a texture/style feature z_s to capture global semantic information (e.g. gender, color, hairstyle, and eyeglasses). Similar approaches have been tried in Diffusion Models in [21, 29], but these techniques are limited to learning a single controllable latent space. Recently, several studies [30, 42] have proposed to have multiple external disentangled spaces of Diffusion Models. However, the image generation performance or the feature disentangling performance are limited.

In this paper, we propose a novel framework FDiff as shown in Fig. 1 to effectively learn two latent spaces to enhance the controllability of diffusion models. Inspired by [20, 28], we add a *Content Encoder* that learns a spatial layout mask and a *Style Encoder* that outputs a flattened semantic feature to condition the diffusion model during training (Section 3.1). The content and style features are injected separately into the UNet [33] to ensure that they encode different semantic factors of an image.

Although decomposing content and style information from an image extends the dimension of the control space, strictly assuming conditional independence between the features during the sampling (i.e., $z_c \perp\!\!\!\perp z_s|x$) may not always be ideal. For example, *face structure* (e.g. square or round face) encoded in the content feature and *gender* (e.g. male or female) encoded in the style feature [28] are not independent given an image. This is because those two features in reality are correlated in a way that males tend to have square jaws with more protruded chins compared to those of females [4]. However, an existing sampling method Composable Diffusion Models (CDM) [24] assumes that conditioning inputs are independent and hence shows unnatural compositions for certain prompts like ‘a flower’ and ‘a bird’ (e.g., Fig.6 in [24] and Fig. 6 in this paper).

To this end, we extend the formulation in [24] and propose *Generalized Composable Diffusion Models* (GCDM) to support compositions during inference when the conditional inputs are not necessarily independent (Section 3.2). Experiment results consistently show that GCDM has better performance in compositing multiple conditions in terms of realism than CDM by breaking the




Fig. 1: Overview of FDiff. Details are provided in Section 3.1.

conditional independence assumption. For instance, from the first row of Fig. 6, we can see that the unnatural composition of ‘a yellow flower’ and ‘a red bird’ gets improved by considering the dependence between the conditions and thus GCDM can naturally place the objects together.

To improve the results further, we leverage the inductive bias [1, 7, 8] of Diffusion Models that learn low-frequency layout information in the later timesteps and high-frequency or imperceptible details in the earlier timesteps. We use a predefined controllable timestep-dependent weight schedule to compose the content and style codes during generation.

2 Preliminaries and Related Works

Preliminaries and related works on Diffusion Models [36] such as DDPM [14] and LDM [32] are provided in Section D in the Supplementary.

Guidance in Diffusion Models:

Some recent works have explored modeling the conditional density $p(x_t|c)$ for controllability. Dhariwal et al. [10] proposed to use a pretrained classifier but finetuning a classifier that estimates gradients from noisy images, which increases the complexity of the overall process [16]. Ho et al. [16] proposed to use an implicit classifier while Composable Diffusion Models [24] (CDM) extend the classifier-free guidance approach to work with multiple conditions assuming conditional independence.

Conditional Diffusion Models:

Conditional Diffusion Models have been explored in diverse applications showing state-of-the-art performance in text-to-image generation (DALLE2 [31], Imagen [35], Parti [43]). These methods use pretrained embeddings (e.g., CLIP) that support interpolation but not further editability. Instructpix2pix [2] proposed to generate synthetic paired data via pretrained GPT-3 [3] and StableDiffusion, with which conditional Diffusion Models are trained. DiffAE [29] proposed to learn a semantic space that has nice properties making it suitable for image manipulation. However, a single latent space capturing all the information makes it difficult to isolate attributes to manipulate. Recently, ControlNet and T2Iadapter [26, 44] showed impressive performance in conditioning image generation. They use additional auxiliary networks and layers that are trained to encode structure into pretrained text2image Diffusion Models. However, our architecture is particular to reference-based image translation, the proposed GCDM and timestep scheduling are generic and applicable to any multi-conditioned Diffusion Models beyond image translation.

Disentangling Conditional Diffusion Models:

DisDiff [42] is one of the recent closely related works. With pretrained Diffusion Models, they train an additional encoder and a decoder to have disentangled representations. However, their results are limited to low resolution (64) and the disentanglement is less convincing than ours (e.g., Fig. 3 and 4 in their paper). DEADiff [30] is the most recent related work. They propose to train an additional module of pretrained Diffusion Models to extract the content and




Fig. 2: Conceptual illustration of CDM and GCDM. (a) The result based on CDM can be outside of manifold while the joint guidance stays on manifold. (b) GCDM trades off between the independent guidance by CDM (stronger conditioning, can be unrealistic) and the joint guidance (more realistic).

the style features from the reference image and use them to guide conditional image generation. However, as shown in Section 4, we have observed that their disentangling performance is limited in representing high-level style feature (e.g., facial attributes) and content feature preserving the spatial information.

Inference only Editing in Diffusion Models:

SDEdit [25] enables structure-preserving edits while Prompt-to-prompt [13] modifies the attention maps from cross-attention layers to add, remove, or reweigh the importance of an object in an image. DiffusionCLIP [19], Imagic [18] and Unitune [40] propose optimization-based techniques for text-based image editing. Textual Inversion [11] and DreamBooth [34] finetune pretrained models using a few reference images to get personalized models. Though the above techniques are helpful with editing, most of these methods require computationally expensive optimization, modify the weights of pretrained model for each sample, and/or don't support fine-grained controllability for reference-based image translation. The closest related work to ours is DiffuseIT [21]. They enabled reference and text-guided image translation by leveraging Dino-ViT [5] to encode content and style. However, their approach requires costly optimization during inference and doesn't support controlling the final generation.

3 Proposed Method

Our framework is based on the LDM [32] architecture as it is faster to train and sample from, compared to pixel-based diffusion models. Let x be an input image and E_{LDM} and D_{LDM} be the pretrained and fixed encoder and decoder respectively. The actual input space for our diffusion model is the low-dimensional latent space $x_0 = E_{LDM}(x)$. The output of the reverse diffusion process is the low dimensional latent \hat{x}_0 which is then passed through the pretrained decoder as $\hat{x} = D_{LDM}(\hat{x}_0)$.

3.1 Learning Content and Style Latent spaces of FDiff

Inspired by DiffAE [29] and similar approaches in GANs [20], we introduce a content encoder $E_c(\cdot; \psi)$ and a style encoder $E_s(\cdot; \phi)$ in our framework as shown in Fig. 1. The objective for training FDiff is formulated as:

$$\min_{\theta, \psi, \phi} \mathbb{E}_{x_0, \epsilon_t} [\|\epsilon_t - \epsilon(x_t, t, E_c(x_0; \psi), E_s(x_0; \phi); \theta)\|_2^2], \quad (1)$$

where x_t is from the forward process, i.e., $x_t = q(x_t|x_0)$. To ensure that the encoders capture different semantic factors of an image, we design the shape of z_s and z_c asymmetrically as done in [6, 17, 20, 22, 28, 39]. The content encoder $E_c(x_0; \psi)$ outputs a spatial layout mask $z_c \in \mathbb{R}^{1 \times \frac{h}{4} \times \frac{w}{4}}$ where w and h are the width and height of x_0 latent. On the other hand, $E_s(x_0; \phi)$ outputs $z_s \in \mathbb{R}^{512 \times 1 \times 1}$ after global average pool layer to capture global high-level semantics. At each layer of the denoising UNet $\epsilon(\cdot; \theta)$, the style code z_s is applied using channel-wise affine transformation while the content code z_c is applied spatially. In specific, let \odot denote the element-wise product (i.e., Hadamard product), let \otimes denote outer product, and let $\mathbf{1}$ denote all one's vector/matrix where subscript denotes dimensionality. We define a variant of adaptive group normalization layer (AdaGN) in the UNet as:

$$\begin{aligned} \text{AdaGN}(h^\ell) = & [\mathbf{1}_C \otimes (\mathbf{1}_{H,W} + t_1 \varphi^\ell(z_c))] \odot [(\mathbf{1}_C + \zeta^\ell(z_s)) \otimes \mathbf{1}_{H,W}] \\ & \odot [[t_2 \otimes \mathbf{1}_{H,W}] \odot \text{GN}(h^\ell) + [t_3 \otimes \mathbf{1}_{H,W}]], \end{aligned} \quad (2)$$

where φ^ℓ is the spatial-wise content-specific term operating down or up-sampling at ℓ -th layer to make the dimensions of $\varphi^\ell(z_c)$ and h^ℓ match. ζ^ℓ is the channel-wise style-specific term that is implemented as an MLP layer. h^ℓ is (hidden) input to the ℓ -th layer, and t terms are timestep-specific adjustment terms inspired by Eq. 7 in DiffAE [29]⁴. When C, H, W are all 1 (i.e., the scalar case), the equation can be simplified to reveal the basic structure as $(1 + t_1 \varphi(z_c))(1 + \zeta(z_s))(t_2 \text{GN}(h) + t_3)$. t_2 and t_3 are affine parameters adjusting the input h to the specific timestep. t_1 is a content-specific timestep adjustment term allowing the denoising networks $\epsilon(\cdot; \theta)$ to decide to strengthen/weaken the content effect since the content feature mostly encodes the low-frequency information, which doesn't play an important role in denoising process over the earlier timesteps [1, 7, 8]. The $(1 + _)$ structure reveals the residual architecture of the content and style. It enables the denoising networks $\epsilon(\cdot; \theta)$ to work even when the content or/and the style condition are not given.

3.2 Generalized Composable Diffusion Models (GCDM)

Once the disentangled representations z_c and z_s are obtained, the next thing we need to consider is how to incorporate them into the Diffusion sampling process. Although CDM has provided a fundamental ground for composing multiple conditions, it has an inherent limitation that conditional independence is assumed (i.e., $z_c \perp\!\!\!\perp z_s | x$), which may not always hold in practice.

To this end, we introduce GCDM derived by breaking the conditional independence assumption. The conceptual benefits of GCDM over CDM are illustrated in Fig. 2. (a) shows an example that the content and the style guidances from CDM generate unrealistic samples because the naively combined guidance (in green) could be outside the manifold. On the contrary, the joint

⁴ The proposed AdaGN by DiffAE is formulated as $z_s(t_2 \text{GN}(h) + t_3)$

guidance helps keep the generation within the manifold. (b) visualizes the proposed GCDM which can be seen as a linear interpolation between CDM and the joint guidance. On top of the limited control space of CDM, GCDM can additionally guide the sampling process more realistic.

Deriving GCDM formulation. We omit the parameter notation θ of the denoising networks ϵ for brevity purposes.

Diffusion Models $\epsilon(x_t, t)$ can be seen as estimating the score of data distribution $\nabla_{x_t} \log p(x_t)$ ⁵ [24, 37, 38, 41], from which we can derive the GCDM formulation. We first start from conditional density $p(x_t|z_c, z_s)$ that we want to have during the sampling process. Taking log and partial derivative w.r.t. x_t , we can have

$$\nabla_{x_t} \log p(x_t|z_c, z_s) = \nabla_{x_t} \log p(x_t, z_c, z_s) = \nabla_{x_t} \log p(x_t)p(z_c, z_s|x_t). \quad (3)$$

As opposed to CDM, we do not assume conditional independence, so that the $p(z_c, z_s|x_t)$ term cannot be easily decomposed into $p(z_c|x_t)p(z_s|x_t)$. By breaking the conditional independence assumption, we have

$$\nabla_{x_t} \log p(x_t|z_c, z_s) = \nabla_{x_t} \log p(x_t)p(z_c|x_t)p(z_s|x_t) \left(\frac{p(z_c, z_s|x_t)}{p(z_c|x_t)p(z_s|x_t)} \right). \quad (4)$$

The terms having a form of $\nabla_{x_t} \log p(z|x_t)$ can be seen as implicit classifiers which can be decomposed into $\nabla_{x_t} \log \frac{p(x_t|z)}{p(x_t)}$, where $p(x_t)$ and $p(x_t|z)$ can be modeled by Diffusion Models. By decomposing the implicit classifiers, we have

$$\nabla_{x_t} \log p(x_t|z_c, z_s) = \nabla_{x_t} \log \frac{p(x_t|z_c)p(x_t|z_s)}{p(x_t)} \left(\frac{p(x_t|z_c, z_s)p(x_t)}{p(x_t|z_c)p(x_t|z_s)} \right). \quad (5)$$

Similar to previous studies [15, 24], we can add hyperparameters to control each of the guidance terms and substitute the score function with the denoising networks ϵ , with which we can finally have GCDM formulation in Definition 1. Please see Section B in the Supplementary for the full derivations.

Definition 1 (Generalized Composable Diffusion Models (GCDM)).

The score function of GCDM is the unconditional score function plus a convex combination of joint and independent guidance terms formalized as:

$$\begin{aligned} \nabla_{x_t} \log \tilde{p}_{\alpha, \lambda, \beta_c, \beta_s}(x_t|z_c, z_s) &\triangleq \epsilon(x_t, t) + \alpha \left[\underbrace{\lambda(\epsilon(x_t, t, z_c, z_s) - \epsilon(x_t, t))}_{\nabla_{x_t} \log p(z_c, z_s|x_t)} \right. \\ &\quad \left. + (1 - \lambda) \sum_{i=\{c, s\}} \beta_i \left(\underbrace{\epsilon(x_t, t, z_i) - \epsilon(x_t, t)}_{\nabla_{x_t} \log p(z_i|x_t)} \right) \right], \end{aligned} \quad (6)$$

⁵ It can be also seen as gradient-based MCMC sampling of Energy-Based Models [24]

where $\alpha \geq 0$ controls the strength of conditioning, $\lambda \in [0, 1]$ controls the trade-off between joint and independent conditioning, and β_c and β_s controls the weight for the content and style features respectively under the constraint that $\beta_c + \beta_s = 1$.

Intuitively, we can strengthen the style effect by increasing β_s and vice versa for the content. We can also increase λ if we want to generate more realistic outputs. α can control the overall strength of the guidance terms. Visualizations of the effects of each hyperparameter are provided in Fig. 9 and Section 4.3.

We next show some of the interesting features of GCDM. First, GCDM generalizes simple joint guidance, CDM, and Classifier-Free Guidance [15] (CFG).

Proposition 1 (GCDM Generalizes Joint Guidance, CDM and CFG).

If $\lambda = 1$, then GCDM simplifies to joint guidance:

$$\nabla_{x_t} \log \tilde{p}_{\lambda=1}(x_t|z_c, z_s) = \underbrace{\epsilon(x_t, t) + \alpha(\epsilon(x_t, t, z_c, z_s) - \epsilon(x_t, t))}_{\text{Joint Guidance}} = \nabla_{x_t} \log p(x_t|z_c, z_s). \quad (7)$$

If $\lambda = 0$, then GCDM simplifies to CDM:

$$\nabla_{x_t} \log \tilde{p}_{\lambda=0}(x_t|z_c, z_s) = \epsilon(x_t, t) + \underbrace{\alpha \left[\sum_{i=\{c,s\}} \beta_i (\epsilon(x_t, t, z_i) - \epsilon(x_t, t)) \right]}_{\text{CDM}}. \quad (8)$$

If $\lambda = 0$ and $\beta_s = 0$, then GCDM simplifies to CFG:

$$\nabla_{x_t} \log \tilde{p}_{\lambda=0, \beta_s=0}(x_t|z_c, z_s) = \underbrace{\epsilon(x_t, t) + \alpha \beta_c (\epsilon(x_t, t, z_c) - \epsilon(x_t, t))}_{\text{CFG}}. \quad (9)$$

The proof is simple from inspection of the GCDM definition. Second, the GCDM PDF \tilde{p} from Definition 1 is proportional to a nested geometric average of different conditional distributions.

Corollary 1. *The GCDM distribution \tilde{p} is proportional to nested geometric averages of conditional distributions of x_t :*

$$\tilde{p}_{\alpha, \lambda, \beta_c, \beta_s}(x_t|z_c, z_s) \propto p(x_t)^{(1-\alpha)} \left[p(x_t|z_c, z_s)^\lambda \left(p(x_t|z_c)^{\beta_c} p(x_t|z_s)^{(1-\beta_c)} \right)^{(1-\lambda)} \right]^\alpha. \quad (10)$$

The outermost geometric average is between an unconditional and conditional model. Then inside we have the geometric average of the joint and independent conditional, and finally inside the independent conditional we have a geometric average of the independent conditionals. The derivation is provided in Section C in the Supplementary.

How can GCDM be more realistic than CDM? In theory, both GCDM and CDM formulations are derived from the classifier guidance $\nabla_{x_t} \log p(z_c, z_s|x_t)$ as shown in Eq. (7) and (13) in the Supplementary. The better the classifier, the more realistic the guided samples could be. The implicit classifier $p(z_c, z_s|x_t) \propto \frac{p(x_t|z_c, z_s)}{p(x_t)}$ of CDM only considers the conditional independence case (Fig. 3 (a)) while the one in GCDM considers all of the cases of dependency (Fig. 3 (b)).

Since conditional independence cannot always hold in practice, assuming $\tilde{p}(z_c, z_s|x_t)$ is a true conditional, the implicit classifier of GCDM can

be a better approximation of \tilde{p} , i.e., $\mathcal{D}(\underbrace{p(z_c|x_t)p(z_s|x_t)}_{\text{CDM}}, \tilde{p}) \geq \mathcal{D}(\underbrace{p(z_c, z_s|x)}_{\text{GCDM}}, \tilde{p})$, where \mathcal{D} is an arbitrary divergence between two distributions, and equality holds if $z_c \perp\!\!\!\perp z_s|x$. Thus, generated samples by GCDM can be more realistic.




Fig. 3: Graphical comparisons of (a) CDM and (b) GCDM.

3.3 Timestep Scheduling for Conditioning

It has been observed in [1, 7, 8] that low-frequency information, i.e., coarse features such as pose and facial shape are learned in the earlier timesteps (e.g., $0 < \text{SNR}(t) < 10^{-2}$) while high-frequency information such as fine-grained features and imperceptible details are encoded in later timesteps (e.g., $10^0 < \text{SNR}(t) < 10^4$) in the reverse diffusion process.

Inspired by this, we introduce a weight scheduler for z_c and z_s that determines how much the content and the style conditions are applied to the denoising networks. We use the following schedule:

$$w_c(t) = \frac{1}{1 + \exp(-a(t - b))}, \quad w_s(t) = \frac{1}{1 + \exp(-a(-t + b))}, \quad (11)$$

where a is a coefficient for determining how many timesteps content and style are jointly provided while b indicates the timestep at which $w_s(t) \geq w_c(t)$. In most experiments, we use $a = 0.025$ and $b = 550$. When applying, we use weighted form, denoted $\bar{\varphi}$ and $\bar{\zeta}$, of the style and content functions during sampling. They are defined as $\bar{\varphi}(z_c, t) := w_c(t)\varphi(z_c)$ and $\bar{\zeta}(z_s, t) := w_s(t)\zeta(z_s)$, which respectively replace φ and ζ in Eq. 2. We also tried a simple linear weighting and a constant schedule but observed that the proposed schedule gave consistently better results (Section G.2 in the Supplementary).

We additionally report using timestep scheduling during training. It is an interesting future direction showing better decomposition of content and style (Section G.1 in the Supplementary).

4 Experiments

We comprehensively evaluate the proposed training and sampling methods. Implementation details and experimental setup, such as specific descriptions on the baselines and evaluation measures are provided in Section E in the Supplementary.




Fig. 4: Comparison of the proposed model with baselines on FFHQ dataset.



Fig. 5: Comparison of the proposed model with additional baselines on FFHQ dataset.

4.1 Comparison with Existing Works

Qualitative Results.

Fig. 4 visually shows example generations from different techniques. We observe that DiffAE+SDEdit loses content information while DiffAE+MagicMix generates unnatural images that naively combine the two images. This indicates that a single latent space even with additional techniques such as SDEdit and MagicMix is not suitable for reference-based image translation. Please note that we used the best hyperparameters that they suggested in their papers. DiffuseIT and SAE models maintain more content information but do not transfer enough information from the style image and have no control over the amount of information transferred from style.

One important benefit of our proposed method over the baselines is better controllability. First of all, by manipulating λ , we can control how much joint guidance is applied. In Fig. 4, decreasing λ indirectly increases the effect of style from the style image when $\beta_c = 0$ and $\beta_s = 1$, where β_c and β_s are the weights for each conditional guidance. It is because the smaller λ brings more information from the style guidance (ref. $1 - \lambda$ term in Definition 1). For example, the man on the second row has more wrinkles and a beard as λ decreases. Second, given a

fixed value of λ , we can control the amount of the content and the style guidance by controlling β_c and β_s as shown in Fig. 7.

Regarding Fig. 5, the style representations of DEADiff do not have facial attribute information, and the contents do not preserve detailed spatial information. For example, the color and texture information is transferred from the style while the person’s identity is transferred from the content (e.g., (a), (d), and (e)). The results of StarGAN v2 tend to apply the style strongly and to lose the content identity (e.g., (a), (b), and (c)). Due to the inherent limitations of GANs, they also cannot properly generate out-of-mode samples, such as a hat, a baby, and hands (e.g., (a) and (d)).

Table 1: Quantitative comparison using FID and LPIPS on FFHQ dataset.

	DiffuseIT [21]	SAE [28]	DiffAE + SDEdit [25, 29]	DiffAE + MagicMix [23, 29]	StarGAN v2 [9]	DEADiff [30]	Ours ($\lambda = 0.9$)	Ours ($\lambda = 0.6$)	Ours ($\lambda = 0.3$)
FID	29.99	25.06	26.63	84.55	48.48	76.84	11.99	13.40	15.45
LPIPS	0.47	0.39	0.64	0.41	0.55	0.37	0.34	0.42	0.49

Quantitative Results.

Table 1 shows the quantitative comparison in terms of FID and LPIPS metrics on FFHQ dataset. Please see Section E.1 in the Supplementary for detailed evaluation settings. Ours generate images that are realistic as indicated by the lowest FID scores compared with other models while also performing comparable on diversity as measured by LPIPS. Specifically, DiffAE+SDEdit shows the highest LPIPS but does not show a meaningful translation since it is hard to find any of the content information from the results (e.g., Fig. 4). StarGAN v2 shows better LPIPS than ours. This is because they tend to apply the style stronger than ours while losing the content information. However, StarGAN v2 show worse FID than ours because their results are limited in representing out-of-mode samples as shown in Fig. 5 (a) and (d).

DiffAE+MagicMix shows the worst performance because of its unrealistic generation. SAE and DiffuseIT show lower LPIPS scores than ours, indicating that they transfer relatively little information from the style image onto the generated samples (i.e., less diverse). We can also observe that increasing λ (when $\beta_c = 0$ and $\beta_s = 1$) makes LPIPS worse while improving FID. In other words, the stronger the joint guidance is the more realistic but less diverse the generated samples are. This verifies our assumption in Fig. 2 that the joint component has an effect of pushing the generations into the real manifold.

4.2 Effect of GCDM and Timestep Scheduling

We compare SAE [28] (the best-performing baseline) and ours with CDM and GCDM on AFHQ dataset in Table 2. The joint guidance ($\lambda = 1$) gets the lowest FID indicating that the generations are more realistic as it pulls the guided results to be within the real data manifold. This can be also understood as a

Table 2: FID comparisons between SAE and our model with CDM and GCDM on AFHQ dataset.

	SAE	CDM	GCDM ($\lambda = 0.9$)	GCDM ($\lambda = 1.0$)
FID	9.29	10.57	9.75	8.58
LPIPS	0.45	0.59	0.59	0.57

Table 3: Comparisons between CDM and GCDM in FFHQ. Best method without timestep scheduling is highlighted in bold and with scheduling is highlighted with *.

	w/o schedule		w/ schedule		
	CDM	GCDM	CDM	GCDM ($\beta_c = 1$)	GCDM ($\beta_s = 1$)
FID	21.43	14.46	10.50	10.21*	10.61
LPIPS	0.47	0.51	0.31	0.28	0.33*

**Fig. 6:** GCDM vs CDM for text-to-image generation with Stable Diffusion. We can observe that CDM generates unnatural images (e.g., blending two objects) that may be out of the real manifold while GCDM ensures realistic generations by breaking the conditional independence assumption (e.g., combining two objects in a realistic way)

positive effect of GCDM that breaks the conditional independence assumption of CDM. It worths noting that GCDM can be thought of as interpolating between CDM and the joint guidance since FID for GCDM ($\lambda = 0.9$) is in between the joint and CDM. By comparing LPIPS and FID of the variants of GCDM, we can see that the outputs become less diverse as realism is increased. SAE shows worse performance than ours in terms of both diversity and realism. The qualitative comparisons can be found in Fig. 14 in the Supplementary.

Generalizability of GCDM.

Since GCDM is a general sampling method that can be applied without the disentangled features, we also compare the performance of CDM and GCDM in composing text prompts for text-to-image generation using Stable Diffusion V2[32] in Fig. 6. The phrases before and after ‘and’ are used as the first and the second conditions (which are the text embeddings in this case). The full sentence is used to represent joint conditioning. As shown in Fig. 6, CDM sometimes fails in composing multiple conditions if both conditions contain object information. For example, *the red bird* and *the yellow flower* are unnaturally overlapped in most cases. On the other hand, GCDM consistently shows better compositions in the generated images. This demonstrates that GCDM, a generalized formulation of CDM can show better performance for composing multiple conditions realistically. This is because GCDM additionally considers the joint guidance by breaking the conditional independence assumption of CDM.

Effect of Timestep Scheduling.



Fig. 7: Timestep scheduling improved the results of both CDM and GCDM.



Fig. 8: Visualization of the effect of each guidance term on generation.

		Baseline									
Stronger style	α	1.5	3.0	1.5	1.5	1.5	1.5	1.0	1.5	1.5	1.5
	λ	0.9	0.9	0.5	0.25	0.9	0.9	0.9	0.5	0.25	0.9
Stronger content	β_s	1.0	1.0	1.0	1.0	1.0	1.0	1.0	0.0	0.0	1.0
	a	0.025	0.025	0.025	0.025	0.025	0.1	0.025	0.025	0.025	0.025
	b	550	550	550	550	650	550	550	550	550	450

Fig. 9: Visual guidance on how to set the hyperparameters.

To more carefully analyze the effect of timestep scheduling when combined with GCDM or CDM, we alter the timestep scheduling so that there is at least a 0.1 weight on style or content. Specifically, we change the upper and lower bounds of the sigmoid to be 0.1 and 0.9 in Eq. 11, e.g., $w'_c(t) = 0.8w_c(t) + 0.1$. The results can be seen in Table 3 and Fig. 7. Without timestep scheduling, GCDM shows better performance in both FID (realism) and LPIPS (diversity) than CDM. Combined with timestep scheduling, both CDM and GCDM show meaningful improvements in FID in exchange for losing diversity. This is because timestep scheduling improves content identity preservation. Additionally, timestep scheduling with GCDM variants shows better FID or LPIPS than CDM depending on the strength of guidance terms..

4.3 Choice of Hyperparameters

In Fig. 9, we show the results with various sets of hyperparameters that can be used to control the effect of the content, style, and joint guidance during the sampling. The gray dotted box represents a baseline that we start from, and the rest of the columns show the effects of each hyperparameter. As can be seen in the figure, each hyperparameter can be modified to get desirable results. Please note that hyperparameter searching is *not mandatory* since the general performance

under the baseline setting is good in terms of FID and LPIPS. If a user wants to control more subtly though, here is a tip on how to set the hyperparameters. Increasing α and decreasing λ yields stronger style effects. By increasing α , since it does not affect the direction of the guidance (but affects the norm), we can expect that the visual features already incorporated are strengthened, such as hair color. Decreasing λ though affects the direction by adjusting the guidance components to be incorporated (e.g., Fig. 2 (b)), and thus we can expect that new visual features from the style image are applied, such as hair length. However, since a low λ indicates the joint guidance is used less, it can pull the generation outside of the manifold so the output can be less realistic. To simplify the control space, we recommend to fix a and b to be 0.025 and 550, respectively.

4.4 Analysis and Discussion

Further analysis and results on latent interpolation and KNN are in Section F.1 and F.3 in the Supplementary.

Visualization of Each Guidance Term.

The proposed GCDM has guidance from three terms, the joint conditioning and style and content conditionings separately. Fig. 8 shows a comparison of the effect of these terms. From the content guidance (column 3), it can be seen that the generated animals are not exactly the same as the content image but have the exact same structure and pose. Similarly, when only style guidance is used (column 4), the pose is random while the style such as color and fur corresponds to the style image. From columns 5-6, it can be observed that the GCDM results have more semantic information from the style guidance than the results by simply using the joint guidance.

Note that the joint guidance is obtained by breaking the conditional independence assumption of CDM. Since GCDM considers the joint probability $p(z_s, z_c|x)$ as well, it can naturally incorporate both the content and the style information (as argued in Fig. 2 (b)). This can be also visual evidence of why GCDM shows the better composition of multiple objects than CDM in Fig. 6.

Classifier-based comparisons.

To further understand what kind of attributes are encoded in style and content latent spaces, we use pretrained classifiers to predict the attributes of translated images and compare them with the original style and content images. We sample 2000

Table 4: Classifier-based comparisons in FFHQ.

Probability Att. is Equal (%)	x_c			x_s		
	Gender	Age	Race	Gender	Age	Race
SAE	65.95	62.36	50.40	34.05	26.40	27.91
Ours ($\lambda = 0.9$)	65.14	53.79	53.31	34.86	31.60	28.51
Ours ($\lambda = 0.25$)	26.61	25.94	31.73	73.39	56.77	44.48

random images from the test set to use as x_c and another 2000 as x_s to form 2000 content-style pairs. Next, we acquire the translated output x_o and corresponding pseudo labels y_c , y_s , and y_o by leveraging an off-the-shelf pretrained attribute classifier (EasyFace). In Table 4, we show the probabilities that the

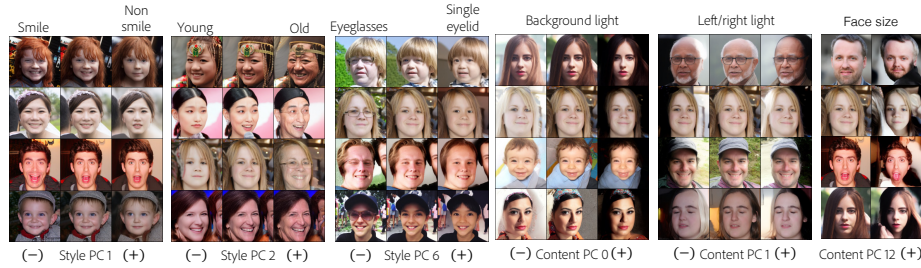


Fig. 10: Visualizations of Principal Components of the style and the content spaces.

final generated image x_o has an attribute from content image as $p(y_c^{att} = y_o^{att})$ and likewise for style image. Both ours and SAE are designed to make z_s encode global high-level semantics, e.g., Gender, Age, etc. Thus, methods would show ideal performance if $y_o^{att} = y_s^{att} \neq y_c^{att}$. We see that most global attributes come from the content image for SAE indicating conservative translations from the style image (as seen in Fig. 4 and lower LPIPS in Table 1). In contrast, ours has a controllable way of deciding the strength of attributes from the style image through λ . Since $\beta_s = 1$ is used, the lower λ yields the stronger style effects (due to $(1 - \lambda)$ term in Definition 1).

Visualizing PCA.

To further visualize the encoded information in the content and the style spaces, we apply PCA algorithm on our pretrained latent spaces, inspired by [12]. Specifically, we first get the style feature and the content feature of all the training images (60000). We then apply PCA to get top 30 Eigenvectors $\mathbf{V}_{\text{style}} = \{v_{\text{style}}^0, \dots, v_{\text{style}}^{29}\}$ and $\mathbf{V}_{\text{content}} = \{v_{\text{content}}^0, \dots, v_{\text{content}}^{29}\}$. The obtained basis vectors are used for shifting each sample. The results are shown in Fig. 10. We can see that the style space mainly encodes high-level semantic information while the content space mainly encodes structural information as we intended.

5 Conclusion

We propose a novel framework FDiff for disentangling the latent space of Diffusion Models. Our content and style encoders trained along with the Diffusion Models and do not require additional objectives or labels to learn to decompose style and content from images. We further propose two sampling methods. The proposed GCDM extends CDM for a more generalized scenario by breaking the conditional independence assumption. It shows better performance when compared with CDM for translation as well as compositing text prompts. We also show that timestep-dependent weight schedules for conditioning inputs can help improve overall results and controllability. Additionally, it is shown that the learned latent spaces have desirable properties from PCA-based and interpolation-based experiments. Quantitative and qualitative evaluation shows the benefits of the proposed sampling techniques.

Acknowledgments. W.C. started this work while on an internship with Adobe. W.C. and D.I. acknowledge partial support from ARL (W911NF-2020-221). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsor.

Bibliography

- [1] Balaji, Y., Nah, S., Huang, X., Vahdat, A., Song, J., Kreis, K., Aittala, M., Aila, T., Laine, S., Catanzaro, B., et al.: ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. arXiv preprint arXiv:2211.01324 (2022)
- [2] Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18392–18402 (2023)
- [3] Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. Advances in neural information processing systems **33**, 1877–1901 (2020)
- [4] Bruce, V., Burton, A.M., Hanna, E., Healey, P., Mason, O., Coombes, A., Fright, R., Linney, A.: Sex discrimination: How do we tell the difference between male and female faces? Perception **22**(2), 131–152 (1993)
- [5] Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the International Conference on Computer Vision (ICCV) (2021)
- [6] Cho, W., Choi, S., Park, D.K., Shin, I., Choo, J.: Image-to-image translation via group-wise deep whitening-and-coloring transformation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10639–10647 (2019)
- [7] Choi, J., Kim, S., Jeong, Y., Gwon, Y., Yoon, S.: Ilvr: Conditioning method for denoising diffusion probabilistic models. arXiv preprint arXiv:2108.02938 (2021)
- [8] Choi, J., Lee, J., Shin, C., Kim, S., Kim, H., Yoon, S.: Perception prioritized training of diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11472–11481 (2022)
- [9] Choi, Y., Uh, Y., Yoo, J., Ha, J.W.: Stargan v2: Diverse image synthesis for multiple domains. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8188–8197 (2020)
- [10] Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Advances in Neural Information Processing Systems **34**, 8780–8794 (2021)
- [11] Gal, R., Alaluf, Y., Atzmon, Y., Patashnik, O., Bermano, A.H., Chechik, G., Cohen-or, D.: An image is worth one word: Personalizing text-to-image generation using textual inversion. In: The Eleventh International Conference on Learning Representations (2023), <https://openreview.net/forum?id=NAQvF08TcyG>
- [12] Härkönen, E., Hertzmann, A., Lehtinen, J., Paris, S.: Ganspace: Discovering interpretable gan controls. Advances in Neural Information Processing Systems **33**, 9841–9850 (2020)

- [13] Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-Or, D.: Prompt-to-prompt image editing with cross attention control. arXiv preprint arXiv:2208.01626 (2022)
- [14] Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
- [15] Ho, J., Salimans, T.: Classifier-free diffusion guidance. In: *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications (2021)*, <https://openreview.net/forum?id=qw8AKxfYbI>
- [16] Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022)
- [17] Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-to-image translation. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 172–189 (2018)
- [18] Kavar, B., Zada, S., Lang, O., Tov, O., Chang, H., Dekel, T., Mosseri, I., Irani, M.: Imagic: Text-based real image editing with diffusion models. arXiv preprint arXiv:2210.09276 (2022)
- [19] Kim, G., Kwon, T., Ye, J.C.: Diffusionclip: Text-guided diffusion models for robust image manipulation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2426–2435 (2022)
- [20] Kwon, G., Ye, J.C.: Diagonal attention and style-based gan for content-style disentanglement in image generation and translation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 13980–13989 (2021)
- [21] Kwon, G., Ye, J.C.: Diffusion-based image translation using disentangled style and content representation. arXiv preprint arXiv:2209.15264 (2022)
- [22] Lee, H.Y., Tseng, H.Y., Huang, J.B., Singh, M., Yang, M.H.: Diverse image-to-image translation via disentangled representations. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 35–51 (2018)
- [23] Liew, J.H., Yan, H., Zhou, D., Feng, J.: Magicmix: Semantic mixing with diffusion models. arXiv preprint arXiv:2210.16056 (2022)
- [24] Liu, N., Li, S., Du, Y., Torralba, A., Tenenbaum, J.B.: Compositional visual generation with composable diffusion models. *Proceedings of the European conference on computer vision (ECCV)* (2022)
- [25] Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: Sdedit: Guided image synthesis and editing with stochastic differential equations. In: *International Conference on Learning Representations* (2021)
- [26] Mou, C., Wang, X., Xie, L., Zhang, J., Qi, Z., Shan, Y., Qie, X.: T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. arXiv preprint arXiv:2302.08453 (2023)
- [27] Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: *International Conference on Machine Learning*. pp. 8162–8171. PMLR (2021)
- [28] Park, T., Zhu, J.Y., Wang, O., Lu, J., Shechtman, E., Efros, A., Zhang, R.: Swapping autoencoder for deep image manipulation. *Advances in Neural Information Processing Systems* **33**, 7198–7211 (2020)

- [29] Preechakul, K., Chatthee, N., Wizadwongsa, S., Suwajanakorn, S.: Diffusion autoencoders: Toward a meaningful and decodable representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10619–10629 (2022)
- [30] Qi, T., Fang, S., Wu, Y., Xie, H., Liu, J., Chen, L., He, Q., Zhang, Y.: Deadiff: An efficient stylization diffusion model with disentangled representations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8693–8702 (2024)
- [31] Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents (2022). <https://doi.org/10.48550/ARXIV.2204.06125>, <https://arxiv.org/abs/2204.06125>
- [32] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10684–10695 (2022)
- [33] Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)
- [34] Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2023)
- [35] Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S.K.S., Ayan, B.K., Mahdavi, S.S., Lopes, R.G., Salimans, T., Ho, J., Fleet, D.J., Norouzi, M.: Photorealistic text-to-image diffusion models with deep language understanding (2022). <https://doi.org/10.48550/ARXIV.2205.11487>, <https://arxiv.org/abs/2205.11487>
- [36] Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: International Conference on Machine Learning. pp. 2256–2265. PMLR (2015)
- [37] Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (2020)
- [38] Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: International Conference on Learning Representations (2021), <https://openreview.net/forum?id=PXTIG12RRHS>
- [39] Tumanyan, N., Bar-Tal, O., Bagon, S., Dekel, T.: Splicing vit features for semantic appearance transfer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10748–10757 (2022)
- [40] Valevski, D., Kalman, M., Matias, Y., Leviathan, Y.: Unitune: Text-driven image editing by fine tuning an image generation model on a single image. arXiv preprint arXiv:2210.09477 (2022)
- [41] Vincent, P.: A connection between score matching and denoising autoencoders. *Neural Computation* **23**(7), 1661–1674 (2011). https://doi.org/10.1162/NECO_a_00142

- [42] Yang, T., Wang, Y., Lu, Y., Zheng, N.: Disdiff: Unsupervised disentanglement of diffusion probabilistic models. NeurIPS (2023)
- [43] Yu, J., Xu, Y., Koh, J.Y., Luong, T., Baid, G., Wang, Z., Vasudevan, V., Ku, A., Yang, Y., Ayan, B.K., Hutchinson, B., Han, W., Parekh, Z., Li, X., Zhang, H., Baldrige, J., Wu, Y.: Scaling autoregressive models for content-rich text-to-image generation (2022). <https://doi.org/10.48550/ARXIV.2206.10789>, <https://arxiv.org/abs/2206.10789>
- [44] Zhang, L., Agrawala, M.: Adding conditional control to text-to-image diffusion models. arXiv preprint arXiv:2302.05543 (2023)