

Versatile Diffusion: Text, Images and Variations All in One Diffusion Model

Xingqian Xu¹, Zhangyang Wang^{2,3}, Eric Zhang¹, Kai Wang¹, Humphrey Shi^{1,3}

¹SHI Labs @ Georgia Tech & UIUC & Oregon, ²UT Austin, ³Picsart AI Research (PAIR)

<https://github.com/SHI-Labs/Versatile-Diffusion>

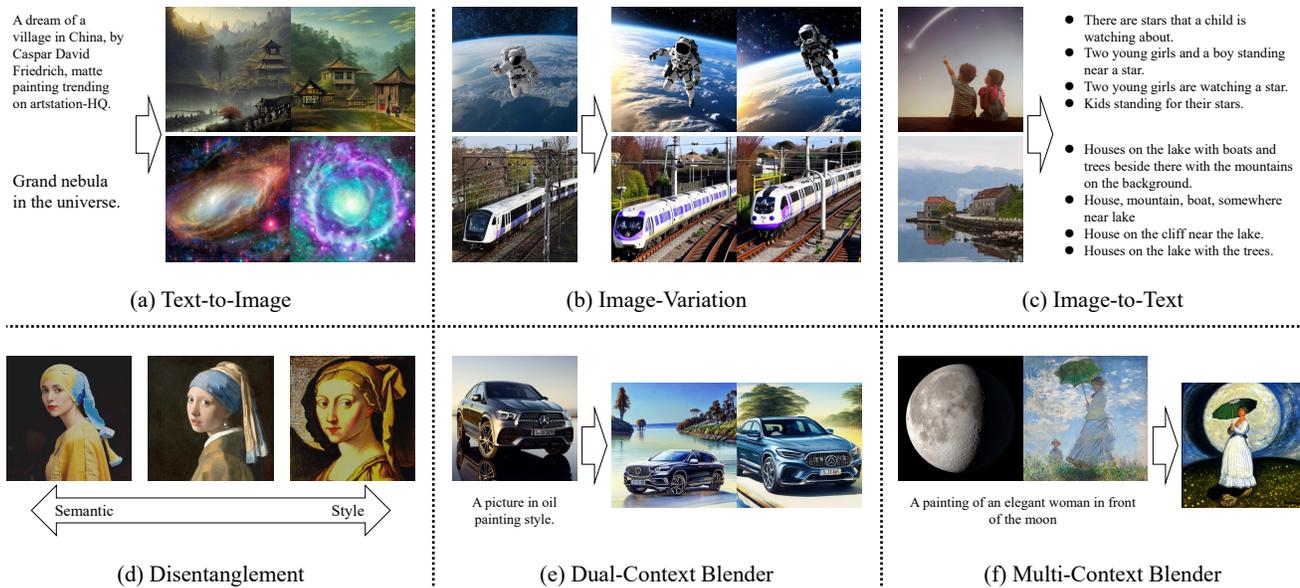


Figure 1: Demo results of our Versatile Diffusion (VD) framework on three out of all primary tasks (*i.e.* Figures a, b, and c) and three derived tasks (*i.e.* Figure d, e, and f). As shown in the captions, the three primary tasks are text-to-image, image-variation, and image-to-text. Figure d demonstrates the disentanglement between image semantics and style. Figure e shows the demo of dual-context blender using one image and one text. Figure f shows the demo of the multi-context blender using multiple images and one text.

Abstract

Recent advances in diffusion models have set an impressive milestone in many generation tasks, and trending works such as DALL-E2, Imagen, and Stable Diffusion have attracted great interest. Despite the rapid landscape changes, recent new approaches focus on extensions and performance rather than capacity, thus requiring separate models for separate tasks. In this work, we expand the existing single-flow diffusion pipeline into a **multi-task multimodal network**, dubbed **Versatile Diffusion (VD)**, that handles multiple flows of text-to-image, image-to-text, and variations in one unified model. The pipeline design of VD instantiates a **unified multi-flow diffusion framework**, consisting of sharable and swappable layer modules that enable the crossmodal generality beyond images and text.

Through extensive experiments, we demonstrate that VD successfully achieves the following: a) VD outperforms the baseline approaches and handles all its base tasks with competitive quality; b) VD enables novel extensions such as disentanglement of style and semantics, dual- and multi-context blending, etc.; c) The success of our multi-flow multimodal framework over images and text may inspire further diffusion-based universal AI research. Our code and models are open-sourced at <https://github.com/SHI-Labs/Versatile-Diffusion>.

1. Introduction

Multi-modality is the “crown jewel” for achieving universal AI. With the attributes of deep learning, methods designed for traditional tasks such as classification, detection, segmentation, *etc.*, have reached near-human level

accuracy. On top of them, multimodal research such as [19, 37, 3, 31] primarily focused on discriminative tasks of jointly recognizing, matching, or understanding multimodal data. Nevertheless, research on multimodal generative models remains scarce. Previously, the best-performing generative vision models, generative adversarial networks (GAN) [34, 7, 35] merely focus on specific domains (*i.e.* faces [35, 10, 94], fonts [99, 45], natural scenes [75, 51], *etc.*); and on specific tasks (inpainting [84, 103, 98], super-resolution [48], image-to-image translation [30, 105], *etc.*).

The recent success of diffusion models [28, 78, 63, 70, 67] has brought new horizons. Diffusion models are likelihood-based models that gradually restore image contents from Gaussian corruptions. It has proved to be effective in bridging modalities and tasks, for instance, unconditional generation [28, 78, 16], density estimation [39], super-resolution [71], and text-to-image generation [56, 63, 70, 67]. The success of diffusion models can be attributed to several aspects. Firstly, their training objectives lead to a more robust training procedure than other approaches like GANs. The iterative refinement inference procedure also expands the model capability at the cost of more running time. Besides, the competitive performance of recent diffusion models such as DALL-E2 [63], Imagen [70], and Stable Diffusion [67] benefits from the remarkable data collection such as LAION [74], CC12M [11], COYO [9], *etc.* The disadvantages of earlier diffusion models, such as the data hunger and high inference costs, are gradually alleviated by more efficient structures and schedulers [78, 43, 73, 29, 67]. Diffusion-based text-to-image methods [63, 70, 67] arguably set new state-of-the-art for multi-modal generative AI. However, those works by far almost exclusively hinge on single-flow diffusion pipelines (illustrated in Section 3); and meanwhile, most of them are trained and evaluated on a single specialized generation task (e.g., text to image) despite being cross-modality.

What is the next move forward, then? We believe in the central role of multimodal, multi-task models in universal AI, and we consider diffusion models to be a promising workhorse to enable so. To fulfill our goal, we proposed *Versatile Diffusion (VD)* that comprehensively solves text, images, and variations within one unified generative model. The *key underlying technique* is a novel multi-flow diffusion framework, that generalizes existing single-flow diffusion pipelines to *handle multiple modalities and tasks simultaneously* while effectively sharing information across them. Thanks to the *larger capacity* as well as capturing crossmodal semantics, VD not only performs well on the aforementioned supported tasks but notably derives many new capabilities including semantic-style disentanglement, cross-modal dual context or multi-context generation (blending), leading to remarkable advances of *empirical performance* for multi-modal generative AI. Our main

contributions are summarized in the following:

- We introduce *Versatile Diffusion (VD)*, a multimodal, multi-task diffusion network that adopts a novel generalized multi-flow pipeline, unlike existing single-flow diffusion models.
- VD solves multiple modalities and tasks in one unified model, including image generation (text-to-image, image-variation), and text generation (image-to-text, text-variation). Through comprehensive experiments, we show that VD outperforms the baselines via scores and quality. For example, VD’s high-quality text-to-image and image-variation results demonstrate that it indeed better captures the context semantics.
- The unique multi-flow multimodal property of VD enables more novel derivative tasks, that may further facilitate downstream users engaged in this technology, including the semantic-style disentanglement, dual-context and multi-context blending, *etc.*

2. Related Works

Multi-modalities are unions of information with different forms, including but not limited to vision, text, audio, *etc.* [83, 4]. Early deep learning work led by Ngiam *et al.* [55] learned a fused representation for audio and video. The similar idea was also adopted across vision and text label [55], and across vision and language [42]. A part of multimodal approaches focused on zero-shot learning, for instance, DiViSE [19] targeted mapping images on semantic space from which unseen category labels can be predicted. Socher *et al.* [76] trained a recognition model with similar ideas in which images were projected on the space of text corpus. [47] shared the same design as DiViSE but was upgraded for a large and noisy dataset. Another set of works [59, 37, 3, 38], focused on increasing classification accuracy via multimodal training: in which [59] and [37] did a simple concatenation on multimodal embeddings; [3] proposed a gated unit to control the multimodal information flow in the network; [38] surveyed FastText [32] with multiple fusion methods on text classification. Meanwhile, multimodal training was also widely adopted in detection and segmentation [22, 25, 31] in one shot. Another topic, VQA [2, 20], conducted cross-modal reasoning that transferred visual concepts into linguistic answers. Methods such as [100, 54] extracted visual concepts into neural symbolics, and [101, 95] learned additional concept structures and hierarchies.

Multimodal generative tasks involve simultaneous representation learning and generation/synthesis [85], in which representation networks [93, 41, 23, 90, 58, 89] with contrastive loss [60, 14, 1, 87, 88] played an essential role. Specifically, our model VD adopts VAEs [41] and CLIP [60] as the latent and context encoders, which are

two critical modules for the network. VD also shares the common cross-modal concepts such as domain transfer [30, 105] and joint representation learning [82, 96, 86].

Diffusion models (DM) [77, 28] consolidate large family of methods including VAEs [41, 90, 65], Markov chains [6, 77, 72, 79], and score matching models [80, 81], *etc.* Differ from GAN-based [23, 7, 35] and flow-based models [66, 40], DM minimizes the lower-bounded likelihoods [28, 80] in backward diffusion passes, rather than exact inverse in flow [66] or conduct adversarial training [23]. Among the recent works, DDPM [28] prompted ϵ -prediction that established a connection between diffusion and score matching models via annealed Langevin dynamics sampling [92, 80]. DDPM also shows promising results on par with GANs in unconditional generation tasks. Another work, DDIM [78], proposed an implicit generative model that yields deterministic samples from latent variables. Compared with DDPM, DDIM reduces the cost of sampling without losing quality. Regarding efficiency, FastDPM [43] investigated continuous diffusion steps and generalized DDPM and DDIM with faster sampling schedules. Another work, [73], replaced the original fixed sampling scheme with a learnable noise estimation that boosted both speed and quality. [29] introduced a hierarchical structure with progressive increasing dimensions that expedite image generations for DM. Regarding quality, [16] compared GANs with DMs with exhaustive experiments and concluded that DMs outperformed GANs on many image generation tasks. Another work, VDM [39], introduced a family of DM models that reaches state-of-the-art performance on density estimation benchmarks. Diffwave [44] and WaveGrad [13] show that DM also works well on audio. [57] improved DDPM with learnable noise scheduling and hybrid objective, achieving even better sampling quality. [53] introduced semantic diffusion guidance to allow image or language-conditioned synthesis with DDPM.

Text-to-image generation, nowadays a joint effort of multimodal and diffusion research, has drawn lots of attention. Among these recent works, GLIDE [56] adopted pretrained language models and the cascaded diffusion structure for text-to-image generation. DALL-E2 [63], a progressive version from DALL-E [64], utilized CLIP model [60] to generate text embedding and adopted the similar hierarchical structure that made 256 text-guided images and then upscaled to 1024. Similarly, Imagen [70] explored multiple text encoders [15, 62, 60] with conditional diffusion models and explores the trade-offs between content alignment and fidelity via various weight samplers. LDM [67] introduced a novel direction in which the model diffuses on VAE latent spaces instead of pixel spaces. Such design reduced the resource needed during inference time, and its latter version, SD, has proven to be equally effective in text-to-image generation.

3. Method

In this section, we will first revisit the fundamentals of diffusion models [77, 28], including the forward-backward processes and training objectives. We will then highlight the multi-flow multimodal framework of Versatile Diffusion (VD), which is a key contribution that makes VD a unified model of multiple tasks. Finally, we will reveal all details of VD, including the choice of VAEs, context encoders, loss functions, *etc.*

3.1. Diffusion basics

The forward diffusion process $p(x_T|x_0)$ is a Markov Chain [28] with T steps that gradually degrade x_0 to x_T with random Gaussian noises (Equation 1).

$$\begin{aligned} q(x_T|x_0) &= \prod_{t=1}^T q(x_t|x_{t-1}) = \prod_{t=1}^T \mathcal{N}(\sqrt{1-\beta_t}x_{t-1}; \beta_t\mathbf{I}) \\ &= \mathcal{N}(\sqrt{\bar{\alpha}_T}x_0; (1-\bar{\alpha}_T\mathbf{I})); \\ \bar{\alpha}_t &= \prod_{s=1}^t \alpha_s; \quad \alpha_t = 1 - \beta_t \end{aligned} \tag{1}$$

Given the forward diffusion process as prior, diffusion models are trained to reverse the process and recover signal x_0 back from x_T by removing the added Gaussian noises. This is known as the backward diffusion process, and each step $p_\theta(x_{t-1}|x_t)$ is sampled from the Gaussian distribution with network predicted mean $\mu_\theta(x_t, t)$ and variance $\Sigma_\theta(x_t, t)$, shown as Equation 2.

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(\mu_\theta(x_t, t), \Sigma_\theta(x_t, t)) \tag{2}$$

The objective function to train a diffusion model is to minimize the variational bound for negative log-likelihood [28] shown in Equation 3. In practice, many works assume deterministic α_t and β_t for step t in Equation 1. Given that both forward and backward processes are Gaussian processes, the objective can then be simplified as the variational weighted l_2 loss between the ground truth and predicted mean.

$$L = \mathbb{E}[-\log p_\theta(x_0)] \leq \mathbb{E} \left[-\log \frac{p_\theta(x_{0:T})}{q(x_{1:T}|x_0)} \right] \tag{3}$$

3.2. Multi-flow multimodal diffusion framework

The core part of Versatile Diffusion (VD) is the multi-flow multimodal diffusion framework capable of generating various forms of outputs (*e.g.* image, text, 3D, *etc.*) conditioned on various crossmodal contexts (*e.g.* image, text, audio *etc.*). A formal definition of a single flow in VD is to synthesize features of modality n using contexts of modality m . One may notice that the well-explored text-to-image

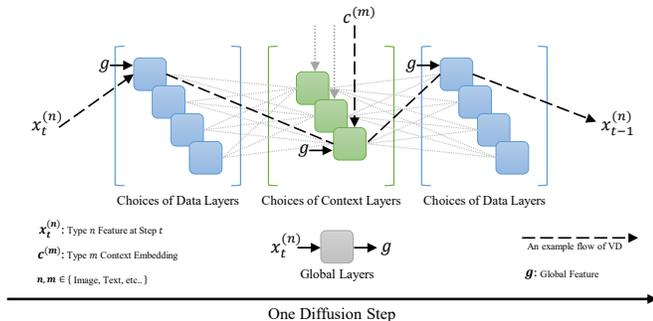


Figure 2: Graphic illustration of one diffusion step of VD’s multi-flow multimodal diffusion framework, in which multiple choices of data layers (blue), context layers (green), and fixed global layers (gray) are involved. The black dash line shows one flow of the model that handles one crossmodal task (*i.e.* text-to-image), in which the top data blocks, the bottom context blocks, and the shared global layers are activated. Other data and context blocks stay silent but will be activated when performing other tasks.

task [21, 63, 70, 67], *i.e.* synthesizing images based on text prompts, matches the definition of a single flow in VD. But the scope of VD goes beyond one single task; particularly in this work, VD is set up to fulfill numerous tasks: text-to-image, image-to-text, and variations, and may further extend to cover more modalities such as 3D, audio, music, *etc.*

Speaking with details, VD handles groups of crossmodal tasks due to its multi-flow framework, in which layers can be activated or muted based on the modalities of the input contexts and output results. As shown in Figure 2, we categorize all diffuser layers into three groups: global layers, data layers, and context layers. The global layers are flow-independent layers that will always be activated. Data layers are output-dependent layers that will be activated when the network generates the corresponding output type. Lastly, context layers are context-dependent layers that will be activated when the corresponding context type is input. Using SD [67] as a reference, the global layers are time-embedding layers; the data layers are residual blocks; and the context layers are cross-attentions. One flow of VD routes the feed-forward pass through the shared global layers and the chosen data and context layers, while other irrelevant layers will stay silent (see Figure 2). Use text-to-image as an example. The t -step intermediate result x_t will be fed to image data blocks and text context blocks to generate the next step result x_{t-1} . Similarly, if our goal is to perform image-variation, we need to use image data blocks and image context blocks.

One may notice that such a multi-flow multimodal framework highly promotes parameter sharing. In this work, our default VD setting is a four-flow model. In order to replicate such four-flow VD, one would require a total of four diffusion models (*i.e.* four times the size of an

SD [67]), while VD reduces the number of parameters by half via its shared layers in the framework. A more generalized version of VD handles $N \times M$ crossmodal tasks with N types of output and M types of context. The size of the model would then become $\mathcal{O}(\max(N, M))$, which is significantly smaller than a vanilla model ensembling that requires an accumulated size of $\mathcal{O}(N \times M)$.

3.3. Versatile Diffusion

Tasks: As mentioned earlier, Versatile Diffusion (VD) is a unified diffusion model for text-to-image, image-to-text, and variations. Text-to-image and image-to-text are two well-known tasks in which the former generates images from text prompts, and the latter generates image captioning. Image-variation (IV) is a fairly new task in which users generate new images that are semantically similar to the reference images. IV differs from SD’s image-to-image (I2I) [67] by two points a) IV diffuses from pure noise while I2I diffuses from images half-mixed with noise; b) IV maintains high-level semantics but relaxes the low-level structures, while I2I only replicates low-level structures and has no guarantee on high-level semantics. Lastly, VD can also generate variations in text due to its multi-flow nature, whose goal is to generate similar expressions from reference text.

Network: The full model of VD includes three components: a) A diffuser that follows our multi-flow multimodal framework described in Sec 3.2; b) VAEs that convert data samples to latent representations; c) Context encoders that encode contexts into embeddings. The overall network diagram is also shown in Figure 3. **Diffuser:** We use the well-adopted UNet [68] with cross attentions [91] as the main structure of our diffuser network. Part of the UNet follows SD [67], where we adopt residual blocks [26] as image data layers and cross-attention as text and image context layers. For text data layers, we propose the fully connected residual blocks (FCResBlock) that expand 768-dimensional text latent vectors into a 320-by-4 hidden feature and follow a similar residual block paradigm with GroupNorms [97], SiLU [18], and skip connections (see Figure 4). **VAE:** We adopt the same Autoencoder-KL [67] like SD as our image VAE. Parallely, we adopt Optimus [49] as our text VAE. Optimus consists of a Bert [15] text encoder and a GPT2 [61] text decoder, by which it can bidirectionally transform sentences into 768-dimensional normally-distributed latent vectors. **Context encoder:** We use both CLIP [60] text and image encoders as VD’s context encoders. Unlike SD, which uses raw text embeddings as context inputs, we use normalized and projected embeddings that minimize the CLIP text-image contrastive loss. In our experiments, we noticed that closer embedding spaces between contexts (*i.e.* image and text) help converge fast and perform better.

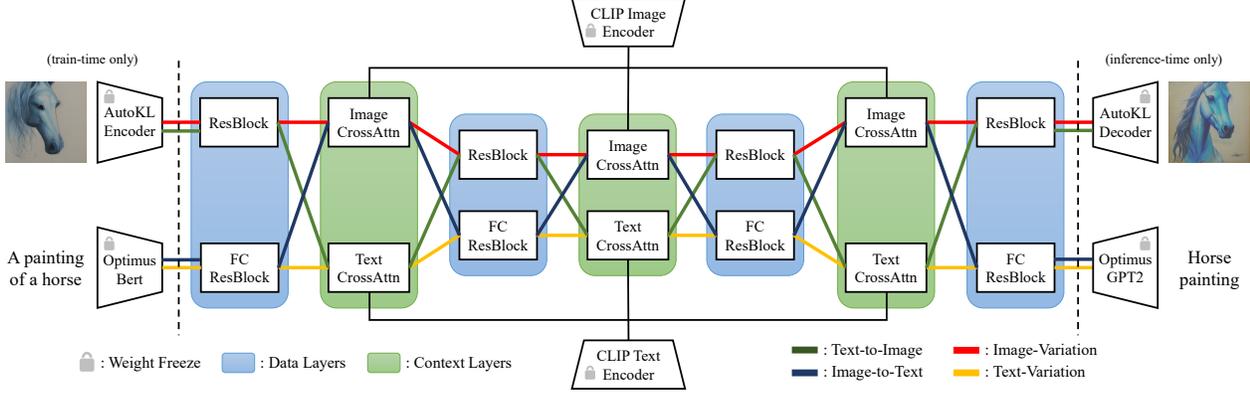


Figure 3: The overall structure of four-flow Versatile Diffusion (VD). Each color line depicts a single flow of VD that represents one supported task (*i.e.* green line for text-to-image). The VAE encoders at the far left are only used in training and are replaced with Gaussian noise inputs during inference. Oppositely, the VAE decoders at the far right are only used in inference for output generation, not train-time loss computation. For simplicity, we hide global layers in this figure. Better viewed in color.

Algorithm 1: Backpropagation of VD

```

 $X = \{x^{(1)} \dots x^{(N)}\}; // N \text{ types data}$ 
 $C = \{c^{(1)} \dots c^{(M)}\}; // M \text{ types context}$ 
 $L_{\theta}(x^{(\cdot)}, c^{(\cdot)}); // \text{Loss with params } \theta$ 
 $\delta_{\theta} = 0; // \text{Param gradients}$ 
for  $x^{(i)} \in X$  do
  for  $c^{(j)} \in C$  do
     $\delta'_{\theta} = \nabla_{\theta} L_{\theta}(x^{(i)}, c^{(j)}); // \text{One flow}$ 
     $\delta_{\theta} = \delta_{\theta} + \delta'_{\theta};$ 
  end
end
Update network with  $\delta_{\theta};$ 

```

Loss: Training VD is surprisingly simple. For each of the flows, we compute the variational weighted l_2 losses described in Equation 3 and do regular backpropagation (see Algorithm 1). Model weights will be updated when the gradients in all flows are accumulated. Besides, when updating the weights, we manually set gradient scales for parameters in data and context layers to better adapt our multi-flow model settings. More information can be found in the Experiments session.

4. Experiments

In this session, we will describe VD’s data and settings, show the performance of VD on primary tasks, and introduce several derived applications empowered by the multi-flow multimodal property of VD.

4.1. Dataset

We used Laion2B-en [74] and COYO-700M [9] as VD’s train data. Both Laion2B and COYO are collections of image-text pairs in English, in which images are collected

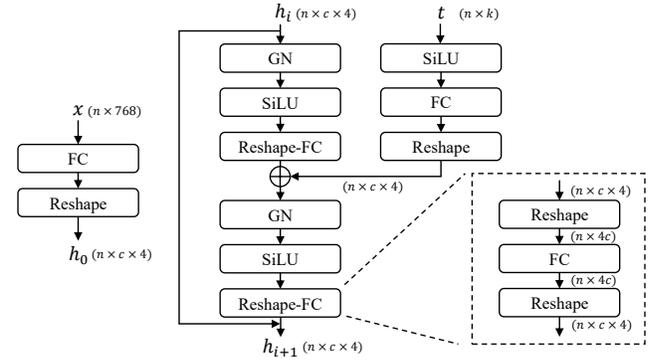


Figure 4: FCResBlock contains two sets of fully connected layers (FC), group normalizations (GN) [97], and sigmoid linear units (SiLU) [18]. x is the input text latent code, t is the input time embedding, and h_i are the intermediate features.

from websites, and the corresponding captions are excerpted from HTML pages. We further filtered all data with the following criteria: a) image-text CLIP similarity scores above 0.3; b) safety scores (*i.e.* NSWF) below 0.3; c) the probability containing watermark below 0.3; d) image aspect ratios within 0.6 to 1.6667; e) image area above $256^2 \times 0.75$. These filtered samples served as the train data for all our VD experiments. Besides, we noticed that the web crawling captions tend to be noisy, so we cleaned them with a customized algorithm described in Supplementary.

4.2. Training

We trained VD progressively with three settings: single-flow, dual-flow, and four-flow, among which the single-flow is an image-variation model; the dual-flow is a text-to-image and image-variation model; and the four-flow is the main VD model with four tasks we majorly described in this work. During training, we kept diffusion settings close

Area of rocks that deep inside the forest, divine domain.

Heavy arms
Gundam penguin mech.

Realistic scenery of Houston Texas city view under a starry sky in hyperrealistic style and ultra HD, 8K.

Red maple on a hill in golden Autumn.



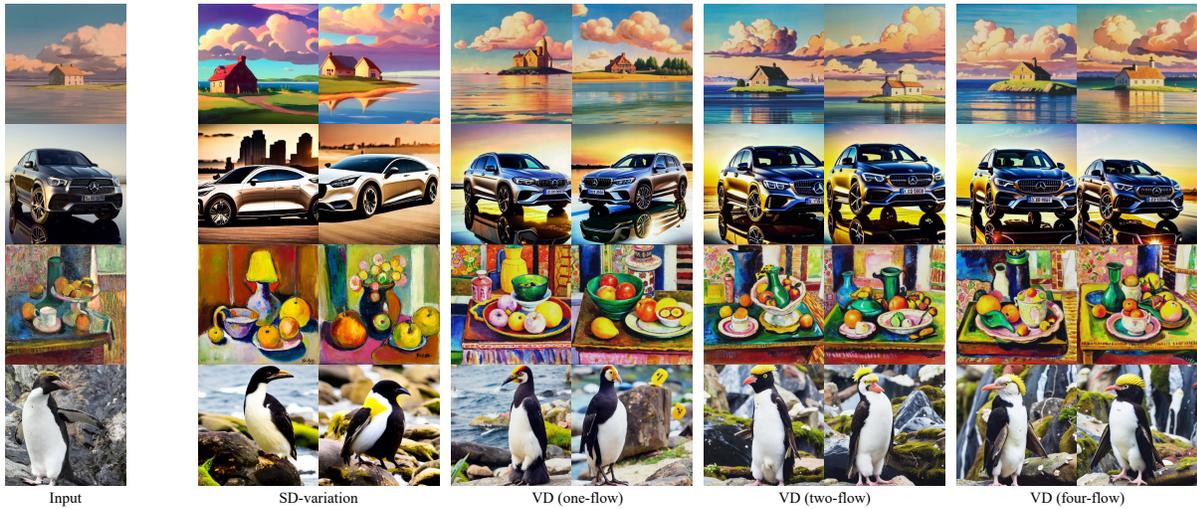
Input

SDv1.4

VD (two-flow)

VD (four-flow)

(a) Text-to-Image performance.



Input

SD-variation

VD (one-flow)

VD (two-flow)

VD (four-flow)

(b) Image-Variation performance.

<p>Input</p>	<p>BLIP</p> <ul style="list-style-type: none"> • Two young kids pointing at the sky. • A girl and boy with a kite flying in the sky. • Two children sitting near each other kids, pointing their fingers at the sun. • Two young kids looking up at the sky. 	<p>VD (four-flow)</p> <ul style="list-style-type: none"> • There are stars that a child is watching about. • Two young girls and a boy standing near a star. • Two young girls are watching a star. • Kids standing for their starts. 	<p>Input</p>	<p>BLIP</p> <ul style="list-style-type: none"> • A house and small body of water. • The mountains behind the water. • A boat docked next to a building. • The shore and the water. 	<p>VD (four-flow)</p> <ul style="list-style-type: none"> • Houses on the lake with boats and trees beside there with the mountains on the background. • House, mountain, boat, somewhere near lake. • House on the cliff near the lake. • Houses on the lake with the trees.
<p>Input</p>	<p>BLIP</p> <ul style="list-style-type: none"> • An old house on the water. • An island with houses in it. • A house in the water. • An island in the sky. 	<p>VD (four-flow)</p> <ul style="list-style-type: none"> • A pictures of a rocky houses on the shore of a lake with a Mountains silhouette on the horizon. • Skyport, a houseboat on a rocky cliff, very near the shore, pictured and the shore. • House on a lake, sky overlooking, and on the shore of the mountain. • A house on a lake by the sky. 	<p>Input</p>	<p>BLIP</p> <ul style="list-style-type: none"> • An image of a person in an art gallery. • An abstract painting that looks like it's on the wall. • An oil painting on canvas. • A painting that looks like an image of the scream. 	<p>VD (four-flow)</p> <ul style="list-style-type: none"> • That is why the man when soul steal happens, is alive, if he see the devil? • A terribly oil painting of his life's pain, on a st. man painting " warbl. • A woman standing on a crying city sidewalk painting. • The man is dead by fate afraid.
<p>Input</p>	<p>BLIP</p> <ul style="list-style-type: none"> • Fireworks and some people watching. • Fireworks with lots of colored bursts. • Fireworks and some clouds. • Fireworks over the skyline and mountains. 	<p>VD (four-flow)</p> <ul style="list-style-type: none"> • A fireworks at night. • Fireworks night. • A fireworks after huge fireworks. • Fireworks bright lights at a party in downtown. 	<p>Input</p>	<p>BLIP</p> <ul style="list-style-type: none"> • The night sky, with some milky. • A person walking across a wooden bridge at night. • A boy in front of the milky. • Someone looking at a milky. 	<p>VD (four-flow)</p> <ul style="list-style-type: none"> • A man standing at a tree outdoors. • Man standing on a tree. • Man standing on tree near city on night. • A man taking a camping hike outdoors.

(c) Image-to-Text performance.

Figure 5: These figures show the qualitative comparison between our VD models and prior works, from which we conclude that VD performs well on all three tasks. In text-to-image and image-variation, VD captures semantics from the input context more accurately. In image-to-text, VD generates more creative sentences and has a better chance to describe images with more details.

to DDPM [28] and SD [67], *i.e.*, 1000 diffusion steps and linearly increasing β from $8.5e - 5$ to $1.2e - 2$ according to steps. The learning rates were set to $1.e - 4$ for single-flow and dual-flow, and were set to $5.e - 5$ for four-flow. The single-flow model used SD checkpoint v1.4 [67] as its initial weights, and others continued finetuning the latest checkpoint from the previous models. During training, we set different gradient scales for different layers to best cooperate with the initial weights. One can find these details in Table 1. The effective batch size was 2048 for single-flow, 1024 for dual-flow, and 512 for four-flow. The logic behind the learning rates, batch sizes, and gradient scales is to roughly balance each gradient step while training. All models were trained with 30 million samples on resolution 256, followed by 6.4 million samples on resolution 512. Compared with SDv1.4, which was trained on 500 plus 230 million samples on resolutions 256 and 512, VD’s training cost is more affordable, benefiting researchers in the long run.

	Data(I)	Data(T)	Ctx(I)	Ctx(T)	Global
VD (1-flow)	0.1	–	1.0	–	0.1
VD (2-flow)	0.1	–	1.0	1.0	0.1
VD (4-flow)	0.2	1.0	1.0	1.0	0.1

Table 1: This table shows the gradient scales used by different layers when training various settings of VD. Data(I) means the image data layer, so on and so forth.

4.3. Performance

To the best of our knowledge, VD is the first image-text multi-flow multimodal model that can be evaluated across different tasks. Thus, we chose single-task-focused prior works as our baselines when comparing the performance. Explicitly speaking: we chose SDv1.4 [67] as our text-to-image baseline; SD-variation [33] (*i.e.* a finetuned SD for image-variation) as our image-variation baseline; and BLIP [50] as our image-to-text baseline. We conducted both qualitative and quantitative comparisons between baselines and various versions of VD, *i.e.*, dual-flow and four-flow for text-to-image, and all three models for image-variation. Although DALLE2 [63] and Imagen [70] also achieved SOTA on text-to-image, they were not compared because of no publicly available code and model. For image-to-text (*i.e.* image captioning), we only compare BLIP [50] with our four-flow VD since other settings do not support this task.

Figure 5 compares VD’s qualitative performance with its baseline, in which images in each row are created with the same random seeds for better quality checks. We also compute text-to-image and image-variation FID scores by comparing 30000 randomly generated samples with the validation set of COCO-caption [52]. In Figure 6, we list VD’s performance along with other related works. We

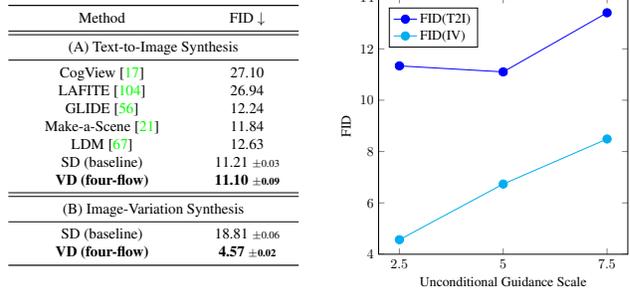


Figure 6: FID scores of VD comparing with baseline and prior approaches, and under various unconditional (classifier-free) guidance scales.

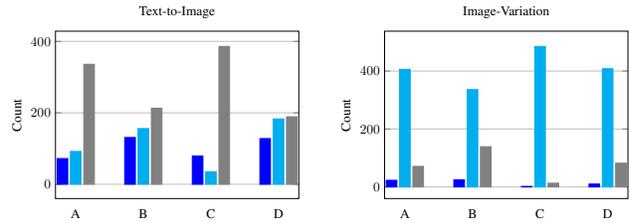


Figure 7: User studies on text-to-image and image-variation in which we count the votes from 4 individual moderators on SD (blue), VD (cyan), or equally good (gray).

also plot the changes in VD’s FID according to the unconditional guidance scale (*i.e.* the classifier-free guidance scale). Lastly, we carried out user studies on 2000 samples from COCO-Caption [52] split by four moderators, in which moderators were asked to vote for better quality or “equally good” (see Figure 7).

Through all results, we not only demonstrated that VD outperforms its baseline on these primary tasks, but reveals the effectiveness of our multi-flow multimodal diffusion framework in which context and data with distinct modalities can be analyzed and generated in one unified model.

4.4. Disentanglement of style and semantic

One exciting discovery of our VD is that it can enhance or reduce image styles from semantics without further supervision. Such a phenomenon inspires us to explore a novel area where disentanglement between styles and semantics can happen on images with arbitrary contents in arbitrary styles. Recall that prior works such as [5, 24] explored similar properties in GAN latent spaces, but their domain of study was restricted to well-aligned data such as faces or churches. To our best knowledge, we are the first group exploring: a) unsupervised semantic and style disentanglement on natural images without domain specifications; b) semantic and style disentanglement on diffusion models’ latent space.

Figure 8 shows the disentanglement results of VD. In practice, we notice that both two-flow and four-flow mod-

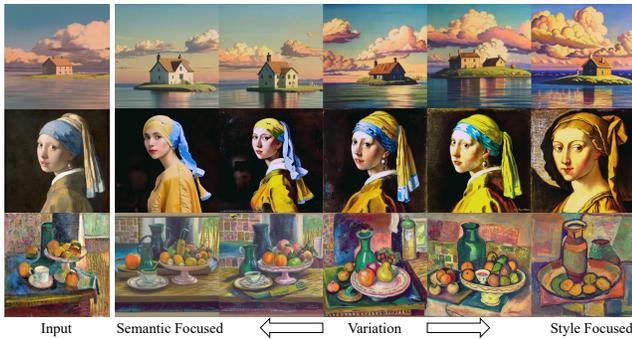


Figure 8: Our VD can disentangle image semantics from styles and vice versa. In this figure, we first generate variations of the input images and then manipulate them focused on either semantics (to the left) or styles (to the right).

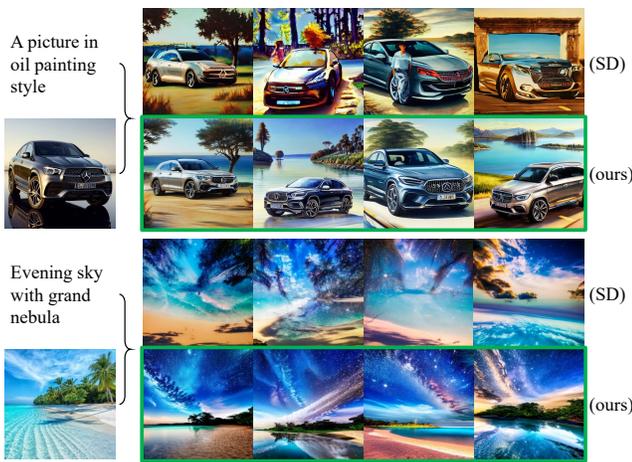


Figure 9: This figure shows images generated from dual-context blender (one image and one prompt). Images without borders are baseline results generated by ensembling SDv1.4 [67] with SD-variation [33]. Images with green borders are VD’s outputs (ours) with a deeper level of mixing. To fairly compare the performance, samples in the same columns use the same random seed and initial noise inputs.

els serve similar performance, while single-flow has slightly lower performance. This may be due to the caption-agnostic and insufficient training that reduced the model’s capacity. More details and analysis can be found in Supplementary.

4.5. Dual- and multi-context blender

Since VD is a unified model for multiple tasks, generation from multi-context becomes a natural extension for VD. Recall that a baseline multi-context generation can be achieved by mixing up diffusion steps from distinct models [53]. However, in practice, we notice such a baseline cannot reach satisfactory results despite doubling the model usage. Figure 9 compares the dual-context results using one text and one image, in which we use the mixing of

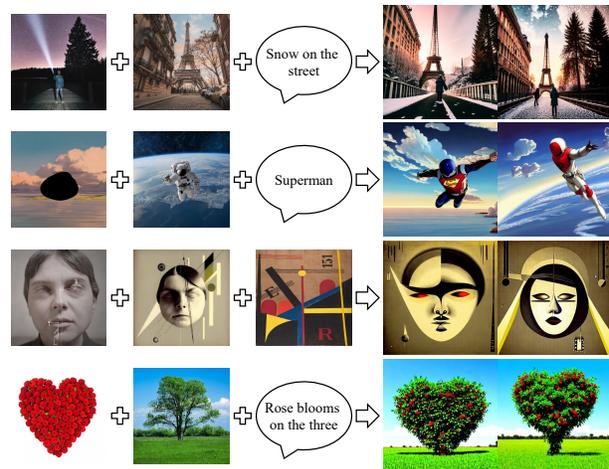


Figure 10: This figure shows images created with VD’s multi-context blender in which multiple images with optional text and masks are applied as contexts. One can notice that VD can smoothly transfer and reconstruct semantic from contexts to outputs.

SDv1.4 [67] (text-to-image) and SD-variation [33] (image-variation) as our baseline (labeled as SD). One may easily notice that VD generates more natural-looking results with fewer distortions. We believe that the good performance of VD is largely attributed to its multi-flow structure, through which intermediate features generated from different contexts can be merged on a much deeper level (*i.e.* layer-level or attention-level), instead of merged on the shallow model-level between diffusion steps. More details regarding mixing levels can be found in Supplementary.

We further expand this task to a more generalized form with multi-context, resulting in the multi-context blender application. The multi-context blender for VD supports an optional text context, several image contexts, and optional image masks in order to guide the generation process with more detail controls. Figure 10 shows the performance of our multi-context blender. Notice that there are other recent works such as [27, 8, 102, 12, 69, 46, 36] focused on the broader image editing topic. We encourage readers to check our Supplementary for more details and comparisons.

5. Conclusion

In this article, we proposed Versatile Diffusion that handles text, image, and variations all in one, from which we generalized a multi-flow multimodal framework that can further extend to new tasks and domains. Through inclusive experiments, we demonstrate that such a multi-flow multimodal diffusion method can perform well on both primary tasks and applications. Moreover, VD can be a heuristic step toward universal AI research.

Acknowledgments. This material is based upon work supported by the National Science Foundation CAREER Award #2239840, and the National AI Institute for Exceptional Education (Award #2229873) by National Science Foundation and the Institute of Education Sciences, U.S. Department of Education. We sincerely thank Professors Thomas S. Huang, Wen-mei Hwu, Zhi-Pei Liang, Mark Hasegawa-Johnson, Trevor Darrell and Bruce Hajek, and Dr. Jie Yang for their constant support and inspiration. We also thank the generous support to the IFP Group and SHI Labs from the Beckman Institute for Advanced Science and Technology & ECE Department & CSL at the University of Illinois at Urbana-Champaign, CS at the University of Oregon, Picsart AI Research (PAIR), various industry gifts (Yamaha, NEC, Adobe, JUMP, Google, IBM, Facebook, SNAP, Kodak, and HP), the Georgia Institute of Technology, and Versatility AI.

References

- [1] Jean-Baptiste Alayrac, Adria Recasens, Rosalia Schneider, Relja Arandjelović, Jason Ramapuram, Jeffrey De Fauw, Lucas Smaira, Sander Dieleman, and Andrew Zisserman. Self-supervised multimodal versatile networks. *Advances in Neural Information Processing Systems*, 33:25–37, 2020. **2**
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015. **2**
- [3] John Arevalo, Tamar Solorio, Manuel Montes-y Gómez, and Fabio A González. Gated multimodal units for information fusion. *arXiv preprint arXiv:1702.01992*, 2017. **2**
- [4] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE transactions on pattern analysis and machine intelligence*, 41(2):423–443, 2018. **2**
- [5] David Bau, Jun-Yan Zhu, Hendrik Strobelt, Bolei Zhou, Joshua B. Tenenbaum, William T. Freeman, and Antonio Torralba. Gan dissection: Visualizing and understanding generative adversarial networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019. **7**
- [6] Yoshua Bengio, Eric Laufer, Guillaume Alain, and Jason Yosinski. Deep generative stochastic networks trainable by backprop. In *International Conference on Machine Learning*, pages 226–234. PMLR, 2014. **3**
- [7] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. **2, 3**
- [8] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. *arXiv preprint arXiv:2211.09800*, 2022. **8**
- [9] Minwoo Byeon, Beomhee Park, Haechon Kim, Sungjun Lee, Woonhyuk Baek, and Saehoon Kim. Coyo-700m: Image-text pair dataset. <https://github.com/kakaobrain/coyo-dataset>, 2022. **2, 5**
- [10] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5799–5809, 2021. **2**
- [11] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3558–3568, 2021. **2**
- [12] Hila Chefer, Yuval Alaluf, Yael Vinker, Lior Wolf, and Daniel Cohen-Or. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *arXiv preprint arXiv:2301.13826*, 2023. **8**
- [13] Nanxin Chen, Yu Zhang, Heiga Zen, Ron J Weiss, Mohammad Norouzi, and William Chan. Wavegrad: Estimating gradients for waveform generation. *arXiv preprint arXiv:2009.00713*, 2020. **3**
- [14] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. **2**
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. **3, 4**
- [16] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. **2, 3**
- [17] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, et al. Cogview: Mastering text-to-image generation via transformers. *Advances in Neural Information Processing Systems*, 34:19822–19835, 2021.
- [18] Stefan Elfving, Eiji Uchibe, and Kenji Doya. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107:3–11, 2018. **4, 5**
- [19] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. *Advances in neural information processing systems*, 26, 2013. **2**
- [20] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. *arXiv preprint arXiv:1606.01847*, 2016. **2**
- [21] Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. Make-a-scene: Scene-based text-to-image generation with human priors. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XV*, pages 89–106. Springer, 2022. **4**
- [22] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the*

- IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. 2
- [23] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 2, 3
- [24] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. *Advances in Neural Information Processing Systems*, 33:9841–9850, 2020. 7
- [25] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 2
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 4
- [27] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 8
- [28] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2, 3, 7
- [29] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *J. Mach. Learn. Res.*, 23:47–1, 2022. 2, 3
- [30] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. 2, 3
- [31] Jitesh Jain, Jiachen Li, MangTik Chiu, Ali Hassani, Nikita Orlov, and Humphrey Shi. Oneformer: One transformer to rule universal image segmentation. *arXiv*, 2022. 2
- [32] Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759*, 2016. 2
- [33] Justin. Experiments with stable diffusion. <https://github.com/justinpinkney/stable-diffusion>. 7, 8
- [34] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 2
- [35] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. 2, 3
- [36] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. *arXiv preprint arXiv:2210.09276*, 2022. 8
- [37] Douwe Kiela and Léon Bottou. Learning image embeddings using convolutional neural networks for improved multi-modal semantics. In *Proceedings of the 2014 Conference on empirical methods in natural language processing (EMNLP)*, pages 36–45, 2014. 2
- [38] Douwe Kiela, Edouard Grave, Armand Joulin, and Tomas Mikolov. Efficient large-scale multi-modal classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. 2
- [39] Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Advances in neural information processing systems*, 34:21696–21707, 2021. 2, 3
- [40] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018. 3
- [41] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2, 3
- [42] Ryan Kiros, Ruslan Salakhutdinov, and Rich Zemel. Multi-modal neural language models. In *International conference on machine learning*, pages 595–603. PMLR, 2014. 2
- [43] Zhifeng Kong and Wei Ping. On fast sampling of diffusion probabilistic models. *arXiv preprint arXiv:2106.00132*, 2021. 2, 3
- [44] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*, 2020. 3
- [45] Praveen Krishnan, Rama Kovvuri, Guan Pang, Boris Vassilev, and Tal Hassner. Textstylebrush: transfer of text aesthetics from a single example. *arXiv preprint arXiv:2106.08385*, 2021. 2
- [46] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. *arXiv preprint arXiv:2212.04488*, 2022. 8
- [47] Angeliki Lazaridou, Elia Bruni, and Marco Baroni. Is this a wampimuk? cross-modal mapping between distributional semantics and the visual world. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1403–1414, 2014. 2
- [48] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017. 2
- [49] Chunyuan Li, Xiang Gao, Yuan Li, Baolin Peng, Xiujun Li, Yizhe Zhang, and Jianfeng Gao. Optimus: Organizing sentences via pre-trained modeling of a latent space. *arXiv preprint arXiv:2004.04092*, 2020. 4
- [50] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. 7

- [51] Chieh Hubert Lin, Hsin-Ying Lee, Yen-Chi Cheng, Sergey Tulyakov, and Ming-Hsuan Yang. Infinitygan: Towards infinite-pixel image synthesis. *arXiv preprint arXiv:2104.03963*, 2021. [2](#)
- [52] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. [7](#)
- [53] Xihui Liu, Dong Huk Park, Samaneh Azadi, Gong Zhang, Arman Chopikyan, Yuxiao Hu, Humphrey Shi, Anna Rohrbach, and Trevor Darrell. More control for free! image synthesis with semantic diffusion guidance. *arXiv preprint arXiv:2112.05744*, 2021. [3](#), [8](#)
- [54] Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B Tenenbaum, and Jiayun Wu. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. *arXiv preprint arXiv:1904.12584*, 2019. [2](#)
- [55] Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. Multimodal deep learning. In *ICML*, 2011. [2](#)
- [56] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. [2](#), [3](#)
- [57] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. [3](#)
- [58] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016. [2](#)
- [59] Deli Pei, Huaping Liu, Yulong Liu, and Fuchun Sun. Unsupervised multimodal feature learning for semantic image segmentation. In *The 2013 International Joint Conference on Neural Networks (IJCNN)*, pages 1–6. IEEE, 2013. [2](#)
- [60] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. [2](#), [3](#), [4](#)
- [61] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. [4](#)
- [62] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67, 2020. [3](#)
- [63] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. [2](#), [3](#), [4](#), [7](#)
- [64] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. [3](#)
- [65] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems*, 32, 2019. [3](#)
- [66] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015. [3](#)
- [67] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. [2](#), [3](#), [4](#), [7](#), [8](#)
- [68] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. [4](#)
- [69] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *arXiv preprint arXiv:2208.12242*, 2022. [8](#)
- [70] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. [2](#), [3](#), [4](#), [7](#)
- [71] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [2](#)
- [72] Tim Salimans, Diederik Kingma, and Max Welling. Markov chain monte carlo and variational inference: Bridging the gap. In *International conference on machine learning*, pages 1218–1226. PMLR, 2015. [3](#)
- [73] Robin San-Roman, Eliya Nachmani, and Lior Wolf. Noise estimation for generative diffusion models. *arXiv preprint arXiv:2104.02600*, 2021. [2](#), [3](#)
- [74] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. [2](#), [5](#)
- [75] Tamar Rott Shaham, Tali Dekel, and Tomer Michaeli. Singan: Learning a generative model from a single natural image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4570–4580, 2019. [2](#)
- [76] Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. Zero-shot learning through cross-modal transfer. *Advances in neural information processing systems*, 26, 2013. [2](#)

- [77] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. 3
- [78] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2, 3
- [79] Jiaming Song, Shengjia Zhao, and Stefano Ermon. A-nice-mc: Adversarial training for mcmc. *Advances in Neural Information Processing Systems*, 30, 2017. 3
- [80] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019. 3
- [81] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020. 3
- [82] Nitish Srivastava and Russ R Salakhutdinov. Multimodal learning with deep boltzmann machines. *Advances in neural information processing systems*, 25, 2012. 3
- [83] Barry E Stein and M Alex Meredith. *The merging of the senses*. The MIT press, 1993. 2
- [84] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2149–2159, 2022. 2
- [85] Masahiro Suzuki and Yutaka Matsuo. A survey of multimodal deep generative models. *Advanced Robotics*, 36(5-6):261–278, 2022. 2
- [86] Masahiro Suzuki, Kotaro Nakayama, and Yutaka Matsuo. Improving bi-directional generation between different modalities with variational autoencoders. *arXiv preprint arXiv:1801.08702*, 2018. 3
- [87] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *European conference on computer vision*, pages 776–794. Springer, 2020. 2
- [88] Yao-Hung Hubert Tsai, Yue Wu, Ruslan Salakhutdinov, and Louis-Philippe Morency. Self-supervised learning from a multi-view perspective. *arXiv preprint arXiv:2006.05576*, 2020. 2
- [89] Aäron Van Den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International conference on machine learning*, pages 1747–1756. PMLR, 2016. 2
- [90] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 2, 3
- [91] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017. 4
- [92] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011. 3
- [93] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, Pierre-Antoine Manzagol, and Léon Bottou. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research*, 11(12), 2010. 2
- [94] Steven Walton, Ali Hassani, Xingqian Xu, Zhangyang Wang, and Humphrey Shi. Stylenat: Giving each head a new perspective. 2022. 2
- [95] Zhonghao Wang, Kai Wang, Mo Yu, Jinjun Xiong, Wenmei Hwu, Mark Hasegawa-Johnson, and Humphrey Shi. Interpretable visual reasoning via induced symbolic space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1878–1887, 2021. 2
- [96] Mike Wu and Noah Goodman. Multimodal generative models for compositional representation learning. *arXiv preprint arXiv:1912.05075*, 2019. 3
- [97] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–19, 2018. 4, 5
- [98] Vahram Tadevosyan Andranik Sargsyan Yadong Mu Xingqian Xu, Shant Navasardyan and Humphrey Shi. Image completion with heterogeneously filtered spectral hints. In *WACV*, 2023. 2
- [99] Shuai Yang, Zhangyang Wang, Zhaowen Wang, Ning Xu, Jiaying Liu, and Zongming Guo. Controllable artistic text style transfer via shape-matching gan. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4442–4451, 2019. 2
- [100] Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Pushmeet Kohli, and Josh Tenenbaum. Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. *Advances in neural information processing systems*, 31, 2018. 2
- [101] Mo Yu, Shiyu Chang, Yang Zhang, and Tommi S Jaakkola. Rethinking cooperative rationalization: Introspective extraction and complement control. *arXiv preprint arXiv:1910.13294*, 2019. 2
- [102] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*, 2023. 8
- [103] Shengyu Zhao, Jonathan Cui, Yilun Sheng, Yue Dong, Xiao Liang, Eric I Chang, and Yan Xu. Large scale image completion via co-modulated generative adversarial networks. *arXiv preprint arXiv:2103.10428*, 2021. 2
- [104] Yufan Zhou, Ruiyi Zhang, Changyou Chen, Chunyuan Li, Chris Tensmeyer, Tong Yu, Jiuxiang Gu, Jinhui Xu, and Tong Sun. Lafite: Towards language-free training for text-to-image generation. *arXiv preprint arXiv:2111.13792*, 2021.
- [105] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 2, 3