EBDM: Exemplar-guided Image Translation with Brownian-bridge Diffusion Models

Eungbean Lee¹, Somi Jeong², and Kwanghoon Sohn^{1,3}

¹ Yonsei University, Seoul, Korea {eungbean,khsohn}@yonsei.ac.kr ² NAVER LABS somi.jeong@naverlabs.com ³ Korea Institute of Science and Technology (KIST), South Korea

Abstract. Exemplar-guided image translation, synthesizing photo-realistic images that conform to both structural control and style exemplars, is attracting attention due to its ability to enhance user control over style manipulation. Previous methodologies have predominantly depended on establishing dense correspondences across cross-domain inputs. Despite these efforts, they incur quadratic memory and computational costs for establishing dense correspondence, resulting in limited versatility and performance degradation. In this paper, we propose a novel approach termed Exemplar-guided Image Translation with Brownian-Bridge Diffusion Models (EBDM). Our method formulates the task as a stochastic Brownian bridge process, a diffusion process with a fixed initial point as structure control and translates into the corresponding photo-realistic image while being conditioned solely on the given exemplar image. To efficiently guide the diffusion process toward the style of exemplar, we delineate three pivotal components: the Global Encoder, the Exemplar Network, and the Exemplar Attention Module to incorporate global and detailed texture information from exemplar images. Leveraging Bridge diffusion, the network can translate images from structure control while exclusively conditioned on the exemplar style, leading to more robust training and inference processes. We illustrate the superiority of our method over competing approaches through comprehensive benchmark evaluations and visual results.

Keywords: Generative model· Image Translation· Diffusion Models · Image Synthesis

1 Introduction

The rising interest in applications of image synthesis has led to a notable surge in demand for image generation capabilities that extend beyond text prompts, emphasizing control through exemplar images or structured inputs. Exemplarguided image translation task aims to generate photo-realistic images conditioned on both a style exemplar image and specific structural controls, such as segmentation masks, edge maps, or pose keypoints.



Fig. 1: Motivation.(a) Existing methods with matching-than-generation framework, (b) Widely used framework based on conditional noise-to-image diffusion model, and (c) our framework based on Brownian bridge diffusion models.

To synthesize images guided by the style of an exemplar and structure controls, pioneer works [13, 31, 42, 43, 48] have emerged. Formulated as an ill-posed problem, these methods globally leverage the style of an exemplar. Despite promising results, these methods overlook local details, which leads to compromised generation quality.

To enhance the capture of local styles from exemplar, significant efforts [20, 39, 54–56, 58, 61] have been explored to establish cross-domain correspondences between input control and exemplar images, thereby imposing local style through a matching process. Zhang *et al.* [58] have explored the construction of cross-domain correspondences using cosine similarity for exemplar-based image translation and subsequent works [54–56, 61] introduced various techniques to reduce computational complexity and address many-to-one matching problems. Moreover, they predominantly capture style information at a coarse scale, which leads to performance degradation, because their methods are significantly influenced by the quality of warped intermediate features derived from sparse correspondences between two domains, often failing to accurately reflect the dynamic nature of matching. This failure leads to local distortion, blurred details, and semantic inconsistency. Furthermore, models leveraging Generative Adversarial Networks (GANs) face the intrinsic limitations of GANs, such as mode collapse, limited diversity, and the out-of-range problem [15, 53].

Recently, diffusion models [8, 34, 40, 41], which generate high-quality images through iterative denoising processes, have attained significant success in the field of image synthesis for their several advantages, including broader distribution coverage, more stable training, and enhanced scalability compared to GANs [2, 8]. As the surge in demand for customized image generation has advanced, Text-to-image (T2I) synthesis, conditioned on text prompts, has been extensively explored in works such as [4, 29, 33]. Beyond mere text prompts, numerous studies have sought solutions to address precise style through model fine-tuning [10, 35], prompt engineering [5, 49]. Additionally, efforts have been made to incorporate structure controls (*i.e.*, edge, depth, mask, pose, *etc.*) [11, 28, 47, 57, 60] as generative guidance.

Although diffusion models have demonstrated impressive performance, exemplarbased image translation remains largely unexplored. First, it is challenging to find an accurate prompt that conveys every desired aspect of an image. Second, it is hard to address exemplar style because fine-tuning offers quality at a high cost, prompt engineering is more affordable but less detailed, and CLIP representations are not sufficient to address all details in visual cues. Lastly, achieving simultaneous conditioning on both style exemplars and structured controls is challenging, particularly because the diffusion process used in such tasks can be highly sensitive to hyperparameters, including the guidance scales of structure control and embeddings.

To solve the above issues, we introduce a sophisticated technique called EBDMs (Exemplar-based image translation with Brownian bridge Diffusion Models) that fully leverages diffusion models. Our method leverages a stochastic Brownian bridge process [19] that directly learns translation between two domains, thus generating images from structure controls without any conditioning mechanism. To explore desired style control, we propose a Global Encoder and Exemplar Network to leverage coarse and fine details from exemplar images. Moreover, the Exemplar Attention Module effectively consilates the texture information from the exemplar into the denoising process. Our method can generate images by conditioning on structure control and style exemplars with single conditioning mask-to-image, edge-to-image [23], and keypoint-to-image [22]. The experimental results demonstrate the superiority of our approach not only performance but also computational efficiency. The contribution of this work can be summarized as follows:

- We introduce the EBDMs, a novel framework leveraging the stochastic Brownian Bridge diffusion process that translates from structure control to a photo-realistic image while effectively exploiting style from exemplars.
- The proposed method formulated the problem into a single-conditioned bridge diffusion process that ensures the training and inference more robust.
- We propose Global Encoder, Exemplar Network, and Exemplar Attention Module to address both global style and detailed texture of exemplar image.
- Extensive experiments demonstrate that our approach achieves favorable performance on various exemplar-guided image translation tasks.

2 Related Works

2.1 Controllable Diffusion Models

Diffusion models [2,8,34,40,41] aim to synthesize images from random Gaussian noise via an iterative denoising process. For customized image generation, recent methods have explored text-guided image generation (T2I) [4,29,33,34,37] and demonstrated extraordinary generative capabilities in modeling the intricacies of complex images. GLIDE [29] aggregated the CLIP texture representations utilizing classifier-free guidance [9]. DALLE-2 [33] proposed a cascade model using the CLIP latent. VQ-Diffusion [4] proposed to learn the diffusion process on the discrete latent space of VQ-VAE [46].

To address the structure controls (such as mask, edge, pose, *etc.*), several works have proposed fine-tuning approach [47] or adaptive models [28, 57, 60] in addition to text prompts. ControlNet [57] proposed an adaptive network to provide structure guidance to T2I models followed by Uni-ControlNet [60] which expands to a unified framework to accept diverse control signals at once. Concurrently, T2I-Adapter [28] introduces a more simple and lightweight adapter. Such methods enable to provide the structural guidance to existing T2I diffusion models thus providing more precise spatial control.

On the other hand, to accurately reflect the style of an exemplar, numerous studies have been conducted such as model fine-tuning [3, 10, 16, 35, 52], prompt engineering [5, 21, 49]. DreamBooth [35] proposed to fine-tune the T2I models with exemplar image and LoRA [10] proposed a more effective tuning method. IP-Adapter [52] proposed decoupled cross-attention to effectively inject exemplar image features into the denoising network. Moreover, Guo *et al.* [5] proposed the image-specific prompt learning method to learn domain-specific prompt vectors. While other methods [6, 26, 27, 36] enable zero-shot editing of an exemplar image based on a target caption. Despite these method capabilities, it is challenging to find the prompt that accurately generates the image a user envisions, mainly because effectively reflects all desired aspects of an image through text, especially those that are difficult or impossible to describe precisely.

2.2 Exemplar-guided Image Translation

The exemplar-guided image translation task involves generating an image based on an input exemplar and structure controls such as an edge, pose, or mask. A major challenge lies in effectively guiding the context within exemplars relative to the input controls. The SPADE [31] framework proposed spatially-adaptive normalization to generate an image from the semantic mask followed by classadaptive [43] and instance-adaptive [42]. While these approaches have shown promising in global-style translation, they overlooked local details compromising generation quality.

To address the local details, significant efforts have been focused on building dense correspondence. Zhang *et al.* [58] proposed building dense correspondence between input semantic and exemplar image. Although their method has shown promising results, their method is limited by many-to-one matching issues and the quadratic computational and memory complexities of dense matching operations, restricting it to capturing only coarse-scale warped features. To alleviate these issues, recent works introduced effective correspondence learning such as GRU-assisted Patch-Match [61], unbalanced optimal transport [54], bilevel feature alignment strategy [55], multi-scale dynamic sparse attention [20], Cross-domain Feature Fusion Transformer [25] and Masked Adaptive Transformer [14]. Although they have demonstrated promising results, their matchingbased framework still suffers from inherent problems such as sparse matching.

Meanwhile, recent progress [11, 39, 50, 51] has leveraged diffusion models to bridge the gap between style exemplars and structural controls. See *et al.* [39] proposed a two-staged framework, which is a matching module followed by a diffusion module. Although they have successfully applied diffusion models, they still heavily rely on a matching-based framework that does not fully utilize the diffusion models. Paint-by-example [50] proposed self-supervised training for image disentanglement and reorganization, while Composer [11] conceptualized an image as a composition of several representations, suggesting a decompose-thenrecompose approach and ImageBrush [51] learns visual instructions. Although they have demonstrated promising results, they offer limited control, typically constrained to structure-preserving appearance changes or uncontrolled imageto-image translation.

3 Preliminaries

3.1 Diffusion Models

The general idea of Denoising Diffusion Probabilistic Model (DDPM) [8] is to generate images from Gaussian noise via T steps of an iterative denoising process. It consists of two processes: the forward process and the reverse process. Given the original data $\boldsymbol{x}_0 \sim q_{data}(\boldsymbol{x}_0)$, the forward diffusion process maps \boldsymbol{x}_0 into noisy latent variables $\{\boldsymbol{x}_t\}_{t=0}^T$ can be obtained as: $\boldsymbol{x}_t = \sqrt{\alpha_t}\boldsymbol{x}_0 + \sqrt{1-\alpha_t}\boldsymbol{\epsilon}$ where $\boldsymbol{\epsilon}$ is the Gaussian noise and $\{\alpha_t\}_{t=0}^T$ is pre-defined schedule. On the other hand, the corresponding reverse process aims to predict the original data \boldsymbol{x}_0 starting from the pure Gaussian noise $\boldsymbol{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ through iterative denoising processes with pre-defined time steps. It is formulated as another Markov chain as $p_{\theta}(\boldsymbol{x}_{t-1} \mid \boldsymbol{x}_t) := \mathcal{N}(\boldsymbol{x}_{t-1}; \boldsymbol{\mu}_{\theta}(\boldsymbol{x}_t, t), \sigma_t^2 \mathbf{I})$ with learned mean and fixed variance. The denoising network $\boldsymbol{\epsilon}_{\theta}$ is trained to predict the noise by minimizing a weighted mean squared error loss, defined as:

$$L(\theta) = \mathbb{E}_{t,\boldsymbol{x}_0,\boldsymbol{\epsilon}}[\|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_{\theta}(\boldsymbol{x}_t, t)\|_2^2].$$
(1)

Similarly, the conditional diffusion models [36, 38] directly inject the condition \boldsymbol{y} into the training objective (eq. 1), such as $L(\theta) = \mathbb{E}_{t,\boldsymbol{x}_0,\boldsymbol{\epsilon}} \|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_{\theta}(\boldsymbol{x}_t,\boldsymbol{y},t)\|_2^2$.

3.2 Brownian Bridge Diffusion Models

A Brownian Bridge Diffusion Model (BBDM) [19] is an image-to-image translation framework based on a stochastic Brownian Bridge process. Unlike DDPM that conclude at Gaussian noise $\boldsymbol{x}_T \sim \mathcal{N}(0, \mathbf{I})$, BBDM assumes that both endpoints of the diffusion process as fixed data points from an arbitrary joint distribution, *i.e.* $(\boldsymbol{x}_T, \boldsymbol{x}_0) \sim q_{\text{data}}(\mathcal{X}, \mathcal{Y})$. The BBDM directly learns image-to-image translation $q(\boldsymbol{x}_0|\boldsymbol{x}_T)$ with boundary distribution $q_{\text{data}}(\boldsymbol{x}_0, \boldsymbol{x}_T)$ independent of any conditional process, that enhances the fidelity and diversity of the generated samples. The forward process of the Brownian Bridge forms a bridge between two fixed endpoints at t = 0 and T:

$$q(\boldsymbol{x}_t \mid \boldsymbol{x}_0, \boldsymbol{y}) = \mathcal{N}(\boldsymbol{x}_t; (1 - m_t) \, \boldsymbol{x}_0 + m_t \boldsymbol{y}, \delta_t \mathbf{I}), \quad \text{where} \quad \boldsymbol{y} = \boldsymbol{x}_T \qquad (2)$$



Fig. 2: Framework overview. The proposed EBDM framework is a based on (a) Brownian Bridge Diffusion Model and composed of (b) Exemplar Network and a (c) Global Encoder. Global Encoder encodes global style information and Exemplar Network extracts texture information from exemplar image. Extracted texture and global information is then used to guide the diffusion process via Exemplar Attention Module and cross-attention, respectively.

where $m_t = t/T$ and variance term $\delta_t = 2(m_t - m_t^2)$.

The reverse process of BBDM aims to predict x_{t-1} given x_t :

$$p_{\theta}(\boldsymbol{x}_{t-1} \mid \boldsymbol{x}_t, \boldsymbol{y}) = \mathcal{N}(\boldsymbol{x}_{t-1}; \boldsymbol{\mu}_{\theta}(\boldsymbol{x}_t, t), \hat{\delta}_t \mathbf{I}),$$
(3)

where δ_t is the variance of Gaussian noise at step t and $\mu_{\theta}(\boldsymbol{x}_t, t)$ is the predicted mean value of the noise, which is network to be learned. The training objective of BBDM is optimizing the Evidence Lower Bound (ELBO), simplified as:

$$\mathbb{E}_{\boldsymbol{x}_{0},\boldsymbol{y},\boldsymbol{\epsilon}}\left[c_{\epsilon t}\|m_{t}(\boldsymbol{y}-\boldsymbol{x}_{0})+\sqrt{\delta_{t}}\boldsymbol{\epsilon}-\boldsymbol{\epsilon}_{\theta}(\boldsymbol{x}_{t},t)\|^{2}\right].$$
(4)

where $c_{\epsilon t}$ is the coefficient term of estimated noise ϵ_{θ} in mean value term, $\tilde{\mu}_t$.

4 Methodology

In this section, we delineate our framework based upon discrete-time stochastic Brownian Bridge diffusion process [19] for Exemplar-guided image translation (Fig. 2). Given a control $I_{\mathcal{X}}$ sampled from domain \mathcal{X} alongside an exemplar image $I_{\mathcal{Y}}$ from domain \mathcal{Y} , the primary objective is to generate a target image $I_{\mathcal{X}\to\mathcal{Y}}$ that retains the structure of $I_{\mathcal{X}}$ embodying the style of $I_{\mathcal{Y}}$. The key to our method is the infusion of style information from $I_{\mathcal{Y}}$ to guide the diffusion trajectory of the target image. To facilitate this, our method integrates three components: a denoising network equipped with an Exemplar Attention Module, a Global Encoder, and Exemplar Network. The Global Encoder extracts global style information of $I_{\mathcal{Y}}$ and the Exemplar Network captures the appearance features of $I_{\mathcal{Y}}$. The Exemplar Attention Module selectively incorporates appearance information into the denoising process.

In the following sections, we present a detailed explanation of the framework (Sec. 4.1), training strategy (Sec. 4.2), and sampling strategy (Sec. 4.3).

4.1 Exemplar-guided Brownian Bridge Diffusion Models

Denoising Network. Employing the Brownian Bridge diffusion process, our denoising U-Net directly learns the translation from the input controls $I_{\mathcal{X}}$ to images $I_{\mathcal{X}\to\mathcal{Y}}$ that preserves the structures of controls. For efficient training and inference, we employ the Stable Diffusion [34] framework. Specifically, given an image I, the encoder \mathcal{E} maps it into a latent space $\mathbf{z} = \mathcal{E}(I)$, subsequently reconstructed by the decoder $\hat{I} = \mathcal{D}(\mathbf{z})$. The denoising U-Net ϵ_{θ} learns to establish the bridge from fixed initial point $\mathbf{x}_T = \mathbf{z}_{\mathcal{X}}$ to the target, $\mathbf{x}_0 = \mathbf{z}_{\mathcal{X}\to\mathcal{Y}}$.

Unlike existing noise-to-image diffusion frameworks [28,57] embed the structural information through intricate frameworks, our approach translates from structural control to images without explicit conditional operation. Consequently, our framework is able to solely focus on exemplar information that fosters enhanced stable training and inference performance.

Global Encoder. The Global Encoder, utilizing DIN0v2 [30], captures the global style information from the exemplar image $I_{\mathcal{Y}}$. Specifically, exemplar image $I_{\mathcal{Y}}$ is processed through the Global Encoder, subsequently, [CLS] token is extracted and passed through a linear layer to encapsulate global style attributes:

$$\tau_{\theta}(\boldsymbol{I}_{\mathcal{Y}}) = \texttt{Linear}\left(\texttt{DINO}(\boldsymbol{I}_{\mathcal{Y}})_{\texttt{[CLS]}}\right) \in \mathbb{R}^{c},\tag{5}$$

where c denotes the dimension of [CLS] token. The global features are utilized as global style information through a cross-attention mechanism, ensuring that the synthesized output accurately reflects the exemplar's global style.

In the context of text-to-image synthesis [8, 34], prior works have extensively leveraged the CLIP image encoder to convey high-level semantic prompts via cross-attention. This approach, however, primarily focuses on the semantic alignment of prompts and images, thereby overlooking the representation of detailed textures. Furthermore, our method does not need textual prompt alignment. Motivated by recent studies [17, 45] that have demonstrated the superior proficiency of DINO [1,30] over CLIP [32] in encapsulating a broader capability of semantic features in images, attributed to its self-supervised learning strategy, our method incorporates the use of a pre-trained DINOv2 encoder to enhance the semantic fidelity of generated images.

Exemplar Network. Notwithstanding the capability of Global Encoder in capturing overarching style information, it is limited to the retention of finegrained details because it encodes exemplar in low resolution (224²). In contrast, the exemplar-guided image translation tasks require higher fidelity to detail. To this end, we introduce Exemplar Network, referred to ψ_{θ} , of which the objective

is to capture the detailed texture information from the exemplar image, thereby compensating for the global information.

The Exemplar Network adopts a siamese configuration akin to a denoising U-Net, streamlined by omitting the redundant layers for enhanced efficiency during training and inference. It encodes the exemplar $z_{\mathcal{Y}}$ into a feature maps $\{F_1^l\}_{l=0}^N$ across N blocks. Additionally, it processes the global information through cross-attention mechanisms in each block. The exemplar features $\{F_1^l\}_{l=0}^N$ are then integrated into the noise prediction branch via Exemplar Attention Module.

Exemplar Attention module. The straightforward approaches to integrate additional features into the denoising network are concatenation [57, 60] or addition [28]/ However, in contrast to existing works in that control features are spatially aligned with the target image, this approach is not suitable for our task because the exemplar image and target control are not spatially aligned. Therefore, we propose an Exemplar Attention Module to integrate the exemplar features from the Exemplar Network, $\mathbf{F}_1^l \in \mathbb{R}^{C \times H \times W}$, into noise prediction features, $\mathbf{F}_2^l \in \mathbb{R}^{C \times H \times W}$ for each l block. First, these features are concatenated into a spatial-wise: $\mathbf{F}_{in}^l = \operatorname{concat}(\mathbf{F}_1^l, \mathbf{F}_2^l) \in \mathbb{R}^{C \times H \times 2W}$. Following this, self-attention is applied to compute the spatial attention across the features:

$$Q = \phi_q^l(\boldsymbol{F}_{in}^l), \quad \boldsymbol{K} = \phi_k^l(\boldsymbol{F}_{in}^l), \quad \boldsymbol{V} = \phi_v^l(\boldsymbol{F}_{in}^l)$$
$$\boldsymbol{F}_{att}^l = \frac{\boldsymbol{Q}\boldsymbol{K}^T}{\sqrt{\boldsymbol{V}}}, \quad \boldsymbol{F}_{EA}^l = \boldsymbol{W}^l \text{Softmax}(\boldsymbol{F}_{att}^l)\boldsymbol{V} + \boldsymbol{F}_{in}^l$$
(6)

where Q, K and V represents query, key and value, respectively, $\phi(\cdot)$ is layerspecific 1×1 convolution operation, and W^l is trainable parameter. Subsequently, exemplar-attended feature $F_{EA}^l \in \mathbb{R}^{C \times H \times 2W}$ is segmented, with portions corresponding to the denoising features are extracted and forwarded toward the output, $F_{out}^l = \text{Chunk}(F_{EA}^l, 2, \dim = 0) \in \mathbb{R}^{C \times H \times W}$. The Exemplar Attention Module computes the region of interest for each query position, a crucial step in effectively directing the denoising steps towards the target exemplar style. This approach enables the denoising process to selectively assimilate features from the Exemplar Network, enhancing the fidelity of the output to the desired stylistic attributes.

Training Objectives. The training process is performed by optimizing the Evidence Lower Bound (ELBO), following BBDM [19], where the marginal distribution is conditioned on $x_{\mathcal{T}}$. Thus, the training objective ELBO in eq. 4 can be simplified as:

$$\mathbb{E}_{\boldsymbol{x}_{0},\boldsymbol{y},\boldsymbol{I}_{\mathcal{Y}},\boldsymbol{\epsilon}}\left[c_{\epsilon t}\left\|m_{t}\left(\boldsymbol{x}_{T}-\boldsymbol{x}_{0}\right)+\sqrt{\delta_{t}}\boldsymbol{\epsilon}-\boldsymbol{\epsilon}_{\theta}\left(\boldsymbol{x}_{t},t,\tau_{\theta}(\boldsymbol{I}_{\mathcal{Y}}),\psi_{\theta}(\boldsymbol{z}_{\mathcal{Y}},\tau_{\theta}(\boldsymbol{I}_{\mathcal{Y}}))\right)\right\|^{2}\right],\tag{7}$$

where $c_{\epsilon t}$ is the loss weighting function that develops into 1/t, and δ_t denotes the preserved variance schedule, $\delta_t = 2(m_t - m_t^2)$.

Algorithm 1 Training

1:	repeat	
2:	$(oldsymbol{x}_T,oldsymbol{x}_0)\sim q_{ ext{data}}(\mathcal{X},\mathcal{Y})$	\triangleright Sample paired data
3:	$oldsymbol{I}_{\mathcal{Y}} \sim q_{ ext{data}}(\mathcal{Y})$	\triangleright Sample exemplar
4:	$t \sim \text{Uniform}(1, \dots, T)$	\triangleright diffusion timesteps
5:	$t_{ref} \leftarrow 0$	\triangleright reference timestep
6:	$oldsymbol{\epsilon} \sim \mathcal{N}(oldsymbol{0},oldsymbol{I})$	\triangleright sample Gaussian noise
7:	$\boldsymbol{G} \leftarrow \tau_{\theta}(\boldsymbol{I}_{\mathcal{Y}})$	\triangleright Forward pass through Global Encoder
8:	$oldsymbol{F} \leftarrow \psi_{ heta}(oldsymbol{x}_{\mathcal{Y}}, t_{ref}, oldsymbol{G})$	\triangleright Forward pass through Exemplar Network
9:	$oldsymbol{x}_t \leftarrow (1-m_t) oldsymbol{x}_0 + m_t oldsymbol{y} + \sqrt{\delta_t} oldsymbol{\epsilon}$	\triangleright Forward bridge diffusion process
10:	$ abla_{ heta} \left\ m_t \left(oldsymbol{y} - oldsymbol{x}_0 ight) + \sqrt{\delta_t} oldsymbol{\epsilon} - oldsymbol{\epsilon}_{ heta} \left(oldsymbol{x}_t, ight) ight.$	$(\boldsymbol{G}, \boldsymbol{F}, t) \ ^2$ \triangleright Gradient descent step
11:	until converged	

Algorithm 2 Sampling

1: $\boldsymbol{x}_T \sim q_{\text{data}}(\boldsymbol{\mathcal{X}})$ ▷ Sample control input 2: $I_{\mathcal{Y}} \sim q_{\text{data}}(\mathcal{Y})$ \triangleright Sample exemplar input 3: $t'_s \leftarrow \{t'_S, \cdots, t'_1\} \sim \{t_T, \cdots, t_1\}$ $\triangleright S$ Inference timesteps 4: $\boldsymbol{G} \leftarrow \tau_{\theta}(\boldsymbol{I}_{\mathcal{Y}})$ ▷ Forward pass through Global Encoder 5: $\boldsymbol{F} \leftarrow \psi_{\theta}(\boldsymbol{x}_{\mathcal{Y}}, \boldsymbol{G})$ ▷ Forward pass through Exemplar Network 6: for s = S, ..., 1 do $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if s > 1, else $\boldsymbol{\epsilon} = 0$ 7: $\boldsymbol{x}_{t'_{s-1}} = c_{xt'_s} \boldsymbol{x}_{t'_s} + c_{yt'_s} \boldsymbol{x}_T - c_{\boldsymbol{\epsilon}t'_s} \boldsymbol{\epsilon}_{\theta} \left(\boldsymbol{x}_{t'_s}, \boldsymbol{G}, \boldsymbol{F}, t'_s \right) + \sqrt{\tilde{\delta}_{t'_s}} \boldsymbol{\epsilon} \triangleright \text{ Take sampling step}$ 8: return x_0

4.2 Training Strategy

The training process is unfolded in two stages. In the first stage, denoising U-Net, which utilizes the Global Encoder and cross-attention mechanism, is trained to integrate the global style cues from the exemplar image. Throughout this phase, the Exemplar Network is not engaged, and pre-trained parameters of VAE and Global Encoder are kept frozen. The primary goal of this stage is to learn the model to translate from the control into high-quality images that simultaneously preserve the structure of the target control and embody the coarse style of the exemplar. This is achieved through a reconstruction manner, wherein the target image is synthesized using its control and the target image itself as the exemplar.

In the second stage, the Exemplar Network and Exemplar Attention Module are incorporated into previously trained denoising U-Net. It enables focused training of the Exemplar Network and the Exempler Attention Module within the denoising U-Net, while the other parameters of the network are kept frozen. The overall training is conducted following the strategy outlined in [58], which employs the predefined exemplar and target pairs. This strategy facilitates a concentrated learning process while emphasizing the detailed integration of the exemplar style and specific characteristics of the target.

Mathad	DeepFashion			CelebA-HQ (Edge)			CelebA-HQ (Mask)		
Method	FID↓	$SWD\downarrow$	$LPIPS\uparrow$	FID↓	$\mathrm{SWD}\!\!\downarrow$	LPIPS↑	FID↓	$\mathrm{SWD}\!\!\downarrow$	L PIPS↑
Pix2PixHD [48]	25.20	16.40	N/A	42.70	33.30	N/A	43.69	34.82	N/A
SPADE [31]	36.20	27.80	0.231	31.50	26.90	0.187	39.17	29.78	0.254
SelectionGAN [44]	38.31	28.21	0.223	34.67	27.34	0.191	42.41	30.32	0.277
SMIS [63]	22.23	23.73	0.240	23.71	22.23	0.201	28.21	24.65	0.301
SEAN [62]	16.28	17.52	0.251	18.88	19.94	0.203	17.66	14.13	0.285
CoCosNet [58]	14.40	17.20	0.272	14.30	15.30	0.208	21.83	12.13	0.292
CoCosNetv2 [61]	12.81	16.53	0.283	12.85	14.62	0.218	20.64	11.21	0.303
UNITE [54]	13.08	16.65	0.278	13.15	14.91	0.213	N/A	N/A	N/A
RABIT [55]	12.58	16.03	0.284	11.67	14.22	0.219	20.44	11.18	0.307
MCL-Net [56]	12.89	16.24	0.286	12.52	14.21	0.216	N/A	N/A	N/A
MIDMs [39]	10.89	10.10	0.279	15.67	12.34	0.224	N/A	N/A	N/A
Ours	10.62	12.40	0.255	11.84	12.10	0.227	12.21	11.34	0.215

Table 3: Quantitative Results in image quality. Comparing our methods with stateof-the-art exemplar-guided image translation methods.

4.3 Sampling Strategy

The inference process is similar to BBDM [19] that employs the deterministic ODE sampler [40]. Given a inference timesteps $\{t'_s\}_{s=1}^S \sim [1:T]$, the sampling process is formulated as:

$$\boldsymbol{x}_{t'_{s-1}} = c_{xt'_s} \boldsymbol{x}_{t'_s} + c_{yt'_s} \boldsymbol{x}_T - c_{\epsilon t'_s} \epsilon_{\theta} \left(\boldsymbol{x}_{t'_s}, \tau_{\theta}(\boldsymbol{I}_{\mathcal{Y}}), \psi_{\theta}(\boldsymbol{x}_{\mathcal{Y}}, \tau_{\theta}(\boldsymbol{z}_{\mathcal{Y}})), t'_s \right) + \sqrt{\tilde{\delta}_{t'_s}} \boldsymbol{\epsilon} \quad (8)$$

where $c_{\epsilon xt}, c_{\epsilon yt}, c_{\epsilon t}$ are weighting coefficients for each terms. The whole training process and sampling process are summarized in Alg. 1 and 2.

5 Experiments

In this section, we present the experimental results of the proposed method. We conduct three tasks to evaluate our model: Edge-to-photo, mask-to-photo, and pose-to-photo. We perform extensive ablation studies to analyze the effect of each essential component of the proposed method. Also, we provide qualitative and quantitative comparisons with state-of-the-art methods. Implementation details and detailed architecture are described in supplementary material. **Datasets.** We conduct three tasks to evaluate our model: Edge-to-photo, mask-to-photo, and pose-to-photo. For mask-guided and edge-guided image generation tasks, the CelebA-HQ [23] dataset is used and we construct the edge maps using the Canny edge detector following [58,61]. For the pose-guided image generation task, we use deepfashion [22] dataset that consists of 52,712 images with a keypoints annotation. For all tasks, the split of train and validation pairs is consistent with CoCosNet [58] policies.

5.1 Qualitative Evaluation

We present a comparison of qualitative results (Fig. 3) with existing methods [20, 54, 58] at three tasks. The results demonstrate that our method effectively transfers the detailed texture from the exemplar to the target, concurrently



Fig. 3: Qualitative Results. Visual comparisons of the proposed EBDM and stateof-the-art methods over three types of exemplar-guided image translation tasks.

Mathad	Dee	pFash	ion	CelebA-HQ (Edge)			
Method	Sem. \uparrow	Col. \uparrow	Tex. \uparrow	Sem. \uparrow	Col. \uparrow	Tex. \uparrow	
Pix2PixHD [48]	0.943	N/A	N/A	0.914	N/A	N/A	
SPADE [31]	0.936	0.943	0.904	0.922	0.955	0.927	
MUNIT [12]	0.910	0.893	0.861	0.848	0.939	0.884	
EGSC-IT [24]	0.942	0.945	0.916	0.915	0.965	0.942	
CoCosNet [58]	0.968	0.982	0.958	0.949	0.977	0.958	
CoCosNet-v2 [61]	0.969	0.974	0.925	0.948	0.975	0.954	
UNITE [54]	0.957	0.973	0.930	0.952	0.966	0.950	
DynaST [20]	0.975	0.974	0.937	0.952	0.980	0.969	
MIDMs [39]	N/A	N/A	N/A	0.915	0.982	0.962	
MATEBIT $[14]$	N/A	N/A	N/A	0.949	0.986	0.966	
Ours	0.932	0.982	0.939	0.920	0.984	0.968	

Table 4: Quantitative metrics of semantic (Sem.), color (Col.), and texture (Tex.) consistency on two datasets compared with state-of-the-art image synthesis methods.

preserving the structure of controls. Notably, in pose-to-photo, our approach exhibits superiority in capturing detailed patterns and minor objects, such as a cap, which other methods often overlook due to the limitation of matching frameworks. These advantages show the capability of our proposed method that fully leverages the diffusion framework that ensures a more holistic and precise depiction. On the other hand, in edge-to-photo and mask-to-photo tasks, while existing methods also achieve photo-realism, they often tend to overfit to the ground truth (*e.g.* UNITE [54]), thereby constraining its generality. However, our method not only accurately transposes the texture of the exemplar but also adeptly conserves the structure. Moreover, the images synthesized through our method demonstrably excel in photo-realistic attributes against other methods.

5.2 Quantitative Evaluation

Evaluation Metrics. We report the Fréchet Inception Distance (FID) [7] and Sliced Wasserstein Distance (SWD) [18] metrics to evaluate the image perceptual quality by reflecting the distance of feature distributions between real images and generated samples. And we also measure LPIPS [59] to evaluate the diversity of translated images. On the other hand, we show the semantic, color, and texture consistency in Tab. 4, also under the same setting as [58].

Image Quality. Tab. 3 presents a quantitative evaluation against state-ofthe-art matching-based methods [20, 39, 54, 58, 61], showing that our method is competitive both on image quality and diversity across various tasks. Additionally, in the mask-to-photo task, our method demonstrates superior performance, whereas matching-based methods struggle due to their reliance on cross-domain matching—a notably arduous endeavor when masks offer scant correspondence cues. Conversely, by leveraging diffusion models, our method iteratively translates images from masks via noise prediction. This enables our approach to excel



Fig. 4: Visual Comparison based on the choice of Exemplar Encoder.



Table 1: Quantitative Results from theAblation Study.



Fig. 5: Qualitative Comparison to Stable

Diffusion based model.

Method	SSIM↑	PSNR↓
SD+ControlNet	0.882	35.30
$\mathbf{SD}{+}\mathbf{ControlNet}{+}\mathbf{CLIP}$	0.894	35.94
Ours	0.901	36.40

 Table 2: Quantitative comparison to Stable diffusion based model.

in scenarios with limited direct correspondences, showcasing its robustness and

adaptability. Consistency. The semantic and style consistency analysis (Tab. 4) evidences

that our method either leads or remains competitive in style relevance scores, encompassing color and texture dimensions. In the pose-to-photo domain, despite achieving scores comparable to other methods [20,58], a visual assessment (Fig. 3) reveals our method's distinct proficiency in retaining intricate details such as patterns or textures. This achievement is attributable to our integrated framework, which combines the Exemplar Network and Global Encoder within a Brownian bridge diffusion model construct. As a result, our methodology not only yields photo-realistic images but also ensures the preservation of texture and style congruence with the exemplar input, underscoring its effectiveness in generating visually coherent outputs.

5.3 Comparison to State-of-the-Arts Diffusion Methods

We compare our framework against prevalent state-of-the-art (SOTA) diffusionbased techniques, as shown in Fig. 5 and Tab. 2. Based on the Stable Diffusion framework [34], we incorporate the ControlNet [57] and IP-Adapter [52] to facilitate structured and stylistic control, respectively. While the existing SOTA method adeptly captures the control structure and generates photo-realistic images, our method more accurately reflects the style of the exemplar. Notably, diffusion-based approaches, conditioned on multiple information including those derived from ControlNet, textual prompts, and image embeddings, tend to be overly sensitive to hyperparameters such as control and embedding guidance

scales. Conversely, our model, predicated on a Brownian Bridge diffusion process and exclusively conditioned on the exemplar, assures a more effective generation process. Moreover, the capacity of the existing methods for transferring the finer details in exemplar is somewhat constrained by their reliance on CLIP embeddings which often overlook small details. In contrast, our framework, underpinned by the Exemplar Network and Exemplar Attention Module, demonstrates superior adeptness in transposing textures from the exemplar.

5.4 Ablation Study

To validate the efficacy of our proposed architecture, we conduct ablation studies focusing on the following configurations: (1) omitting the Global Encoder, (2) utilizing the baseline model [19] integrated with CLIP, (3) implementing DINOv2, and (4) employing our complete architecture on the edge-to-photo translation task. As illustrated in Fig. 4, our findings reveal that the DINOv2-based Global Encoder surpasses the CLIP in generating images with higher detail fidelity. While CLIP effectively captures the general characteristics of the reference image, ensuring a level of resemblance, it does not fully encapsulate the intricacies of the details. Additionally, with our Exemplar Network, inputs with spatial misaligned control and exemplar often result in the generation of "blurry" images when relying exclusively on features of Global Encoder. In contrast, our complete framework demonstrates superior performance across all assessed dimensions, highlighting its architectural advantage. Quantitative assessments further underscore the importance of our design choices, as detailed in Tab. 1.

6 Conclusion

In this study, we presented EBDM, a novel stochastic Brownian bridge diffusionbased approach for exemplar-guided image translation. By leveraging the Brownian Bridge framework, which translates from fixed data points as structural control to photo-realistic images, our method is exclusively conditioned to the style information, thereby the framework more robust and stable. Additionally, we propose the Exemplar Network and Exemplar Attention Module to selectively incorporate the style information from exemplar images into the denoising process. Our method not only stands competitive or surpasses existing methods across the three distinct tasks. Furthermore, our methods also achieve a significant improvement in visual results not only in photorealism but also in the precise transfer of fine details such as patterns and accessories present in the exemplar images.

Acknowledgements

This research was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIP) (NRF2021R1A2C2006703) and the Yonsei Signature Research Cluster Program of 2024 (2024-22-0161).

References

- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: ICCV (2021)
- Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. In: NeurIPS (2021)
- 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. arXiv preprint arXiv:2208.01618 (2022)
- Gu, S., Chen, D., Bao, J., Wen, F., Zhang, B., Chen, D., Yuan, L., Guo, B.: Vector quantized diffusion model for text-to-image synthesis. In: CVPR (2022)
- Guo, J., Wang, C., Wu, Y., Zhang, E., Wang, K., Xu, X., Shi, H., Huang, G., Song, S.: Zero-shot generative model adaptation via image-specific prompt learning. In: CVPR (2023)
- Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-Or, D.: Prompt-to-prompt image editing with cross attention control. In: ICLR (2023)
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: NeurIPS. vol. 30 (2017)
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: NeurIPS (2020)
- 9. Ho, J., Salimans, T.: Classifier-free diffusion guidance. In: NeurIPS Workshop (2022)
- Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. In: ICLR (2022)
- Huang, L., Chen, D., Liu, Y., Shen, Y., Zhao, D., Zhou, J.: Composer: Creative and controllable image synthesis with composable conditions. arXiv preprint arXiv:2302.09778 (2023)
- Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-toimage translation. In: ECCV. pp. 172–189 (2018)
- Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR. pp. 1125–1134 (2017)
- Jiang, C., Gao, F., Ma, B., Lin, Y., Wang, N., Xu, G.: Masked and adaptive transformer for exemplar based image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22418–22427 (2023)
- Kang, K., Kim, S., Cho, S.: Gan inversion for out-of-range images with geometric transformations. In: ICCV (2021)
- Kawar, B., Zada, S., Lang, O., Tov, O., Chang, H., Dekel, T., Mosseri, I., Irani, M.: Imagic: Text-based real image editing with diffusion models. In: CVPR (2023)
- 17. Kwon, G., Ye, J.C.: Diffusion-based image translation using disentangled style and content representation. In: ICLR (2023)
- Lee, C.Y., Batra, T., Baig, M.H., Ulbricht, D.: Sliced wasserstein discrepancy for unsupervised domain adaptation. In: CVPR. pp. 10285–10295 (2019)
- 19. Li, B., Xue, K., Liu, B., Lai, Y.K.: Bbdm: Image-to-image translation with brownian bridge diffusion models. In: CVPR (2023)
- 20. Liu, S., Ye, J., Ren, S., Wang, X.: Dynast: Dynamic sparse transformer for exemplar-guided image generation. In: ECCV. pp. 72–90. Springer (2022)
- 21. Liu, V., Chilton, L.B.: Design guidelines for prompt engineering text-to-image generative models. In: ACM CHI (2022)

- 16 E. Lee et al.
- 22. Liu, Z., Luo, P., Qiu, S., Wang, X., Tang, X.: Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In: CVPR. pp. 1096–1104 (2016)
- Liu, Z., Luo, P., Wang, X., Tang, X.: Large-scale celebfaces attributes (celeba) dataset. Retrieved August 15, 11 (2018)
- 24. Ma, L., Jia, X., Georgoulis, S., Tuytelaars, T., Van Gool, L.: Exemplar guided unsupervised image-to-image translation with semantic consistency. In: ICLR (2019)
- Ma, T., Li, B., Liu, W., Hua, M., Dong, J., Tan, T.: Cfft-gan: cross-domain feature fusion transformer for exemplar-based image translation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 1887–1895 (2023)
- 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: ICLR (2022)
- 27. Mokady, R., Hertz, A., Aberman, K., Pritch, Y., Cohen-Or, D.: Null-text inversion for editing real images using guided diffusion models. In: CVPR (2023)
- Mou, C., Wang, X., Xie, L., Wu, Y., Zhang, J., Qi, Z., Shan, Y., Qie, X.: T2iadapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. arXiv preprint arXiv:2302.08453 (2023)
- Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In: ICML (2022)
- Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. TMLR (2024)
- Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: CVPR. pp. 2337–2346 (2019)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021)
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical textconditional image generation with clip latents. arXiv preprint arXiv:2204.06125 (2022)
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022)
- 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: CVPR (2023)
- Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: ACM SIGGRAPH (2022)
- 37. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic textto-image diffusion models with deep language understanding. In: NeurIPS (2022)
- Sasaki, H., Willcocks, C.G., Breckon, T.P.: Unit-ddpm: Unpaired image translation with denoising diffusion probabilistic models. arXiv preprint arXiv:2104.05358 (2021)
- Seo, J., Lee, G., Cho, S., Lee, J., Kim, S.: Midms: Matching interleaved diffusion models for exemplar-based image translation. In: AAAI (2023)
- 40. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: ICLR (2021)
- Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Scorebased generative modeling through stochastic differential equations. In: ICLR (2021)

- 42. Tan, Z., Chai, M., Chen, D., Liao, J., Chu, Q., Liu, B., Hua, G., Yu, N.: Diverse semantic image synthesis via probability distribution modeling. In: CVPR. pp. 7962–7971 (2021)
- Tan, Z., Chen, D., Chu, Q., Chai, M., Liao, J., He, M., Yuan, L., Hua, G., Yu, N.: Efficient semantic image synthesis via class-adaptive normalization. TPAMI 4 (2021)
- Tang, H., Xu, D., Sebe, N., Wang, Y., Corso, J.J., Yan, Y.: Multi-channel attention selection gan with cascaded semantic guidance for cross-view image translation. In: CVPR. pp. 2417–2426 (2019)
- Tumanyan, N., Bar-Tal, O., Bagon, S., Dekel, T.: Splicing vit features for semantic appearance transfer. In: CVPR. pp. 10748–10757 (2022)
- Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. In: NeurIPS. vol. 30 (2017)
- 47. Wang, T., Zhang, T., Zhang, B., Ouyang, H., Chen, D., Chen, Q., Wen, F.: Pretraining is all you need for image-to-image translation. arXiv preprint arXiv:2205.12952 (2022)
- Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: Highresolution image synthesis and semantic manipulation with conditional gans. In: CVPR. pp. 8798–8807 (2018)
- Witteveen, S., Andrews, M.: Investigating prompt engineering in diffusion models. In: NeurIPS Workshop (2022)
- Yang, B., Gu, S., Zhang, B., Zhang, T., Chen, X., Sun, X., Chen, D., Wen, F.: Paint by example: Exemplar-based image editing with diffusion models. In: CVPR (2023)
- Yang, Y., Peng, H., Shen, Y., Yang, Y., Hu, H., Qiu, L., Koike, H., et al.: Imagebrush: Learning visual in-context instructions for exemplar-based image manipulation. Advances in Neural Information Processing Systems 36 (2024)
- 52. Ye, H., Zhang, J., Liu, S., Han, X., Yang, W.: Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. arXiv preprint arXiv:2308.06721 (2023)
- 53. Yin, F., Zhang, Y., Cun, X., Cao, M., Fan, Y., Wang, X., Bai, Q., Wu, B., Wang, J., Yang, Y.: Styleheat: One-shot high-resolution editable talking face generation via pre-trained stylegan. In: ECCV (2022)
- 54. Zhan, F., Yu, Y., Cui, K., Zhang, G., Lu, S., Pan, J., Zhang, C., Ma, F., Xie, X., Miao, C.: Unbalanced feature transport for exemplar-based image translation. In: CVPR. pp. 15028–15038 (2021)
- 55. Zhan, F., Yu, Y., Wu, R., Cui, K., Xiao, A., Lu, S., Shao, L.: Bi-level feature alignment for versatile image translation and manipulation. In: ECCV (2022)
- Zhan, F., Yu, Y., Wu, R., Zhang, J., Lu, S., Zhang, C.: Marginal contrastive correspondence for guided image generation. In: CVPR. pp. 10663–10672 (2022)
- 57. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: ICCV (2023)
- Zhang, P., Zhang, B., Chen, D., Yuan, L., Wen, F.: Cross-domain correspondence learning for exemplar-based image translation. In: CVPR. pp. 5143–5153 (2020)
- 59. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR. pp. 586–595 (2018)
- Zhao, S., Chen, D., Chen, Y.C., Bao, J., Hao, S., Yuan, L., Wong, K.Y.K.: Unicontrolnet: All-in-one control to text-to-image diffusion models. In: NeurIPS (2024)
- Zhou, X., Zhang, B., Zhang, T., Zhang, P., Bao, J., Chen, D., Zhang, Z., Wen, F.: Cocosnet v2: Full-resolution correspondence learning for image translation. In: CVPR. pp. 11465–11475 (2021)

- 18 E. Lee et al.
- 62. Zhu, P., Abdal, R., Qin, Y., Wonka, P.: Sean: Image synthesis with semantic regionadaptive normalization. In: CVPR. pp. 5104–5113 (2020)
- 63. Zhu, Z., Xu, Z., You, A., Bai, X.: Semantically multi-modal image synthesis. In: CVPR. pp. 5467–5476 (2020)