

MIDMs: Matching Interleaved Diffusion Models for Exemplar-Based Image Translation

Junyoung Seo,^{*1} Gyuseong Lee,^{*1} Seokju Cho,¹ Jiyoung Lee,² Seungryong Kim¹

¹ Korea University, Seoul, Korea

² NAVER AI Lab, Korea

se780@korea.ac.kr, jpl358@korea.ac.kr, seokju_cho@korea.ac.kr, lee.j@navercorp.com, seungryong.kim@korea.ac.kr

Abstract

We present a novel method for exemplar-based image translation, called matching interleaved diffusion models (MIDMs). Most existing methods for this task were formulated as GAN-based matching-then-generation framework. However, in this framework, matching errors induced by the difficulty of semantic matching across cross-domain, e.g., sketch and photo, can be easily propagated to the generation step, which in turn leads to degenerated results. Motivated by the recent success of diffusion models overcoming the shortcomings of GANs, we incorporate the diffusion models to overcome these limitations. Specifically, we formulate a diffusion-based matching-and-generation framework that interleaves cross-domain matching and diffusion steps in the latent space by iteratively feeding the intermediate warp into the noising process and denoising it to generate a translated image. In addition, to improve the reliability of the diffusion process, we design a confidence-aware process using cycle-consistency to consider only confident regions during translation. Experimental results show that our MIDMs generate more plausible images than state-of-the-art methods.

Introduction

Image-to-image translation, aiming to learn a mapping between two different domains, has shown a lot of progress in recent years (Zhu et al. 2017; Isola et al. 2017; Wang et al. 2018; Chen and Koltun 2017; Park et al. 2019). Especially, exemplar-based image translation (Ma et al. 2018; Wang et al. 2019; Zhang et al. 2020; Zhou et al. 2021; Zhan et al. 2022, 2021a) that can generate an image conditioned on an exemplar image has attracted much attention due to its flexibility and controllability. For instance, translating a user-given condition image, e.g., pose keypoints, segmentation maps, or stroke, to a photorealistic image conditioned on an exemplar real image can be used in numerous applications such as semantic image editing or makeup transfer (Zhang et al. 2020; Zhan et al. 2021b).

To solve this task, early pioneering works (Huang et al. 2018; Ma et al. 2018; Wang et al. 2019) attempted to transfer a global style of exemplar. Recently, several works (Zhang et al. 2020; Zhou et al. 2021; Zhan et al. 2022, 2021a)

have succeeded in bringing the local style of exemplar by combining matching networks with Generative Adversarial Networks (GANs) (Goodfellow et al. 2014)-based generation networks, i.e., GANs-based matching-then-generation. Formally, these approaches first establish matching across cross-domain and then synthesize an image based on a warped exemplar. However, the efficacy of such a framework is largely dependent on the quality of warped intermediates, which hinders faithful generations in case unreliable correspondences are established. Furthermore, GANs-based generators inherit the weaknesses of the GAN model, i.e., convergence heavily depends on the choice of hyperparameters (Gulrajani et al. 2017; Arjovsky, Chintala, and Bottou 2017; Salimans et al. 2016; Goodfellow 2016), lower variety, and mode drop in the output distribution (Brock, Donahue, and Simonyan 2018; Miyato et al. 2018).

On the other hand, recently, diffusion models (Sohl-Dickstein et al. 2015; Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2020; Rombach et al. 2021) have attained much attention as an alternative generative model. Compared to GANs, diffusion models can offer desirable qualities, including distribution coverage, a fixed training objective, and scalability (Ho, Jain, and Abbeel 2020; Dhariwal and Nichol 2021; Nichol et al. 2021). Even though the diffusion models have shown appealing performances in image generation and manipulation tasks (Choi et al. 2021; Meng et al. 2021; Kim and Ye 2021), applying this to exemplar-based image translation remains unexplored.

In this paper, we propose to use diffusion models for exemplar-based image translation tasks, called matching interleaved diffusion models (MIDMs), to address the limitations of existing methods (Zhang et al. 2020; Zhou et al. 2021; Zhan et al. 2021a,b, 2022). We for the first time adopt the diffusion models to exemplar-based image translation tasks, but directly adopting this in the matching-then-generation framework similarly to (Zhang et al. 2020) may generate sub-optimal results. To overcome this, we present a diffusion-based matching-and-generation framework that interleaves cross-domain matching and diffusion steps to modify the diffusion trajectory toward a more faithful image translation, as shown in Fig. 1. We allow the recurrent process to be confidence-aware by using the cycle-consistency so that our model can adopt only reliable regions for each iteration of warping. The proposed MIDMs overcome the lim-

^{*}These authors contributed equally.

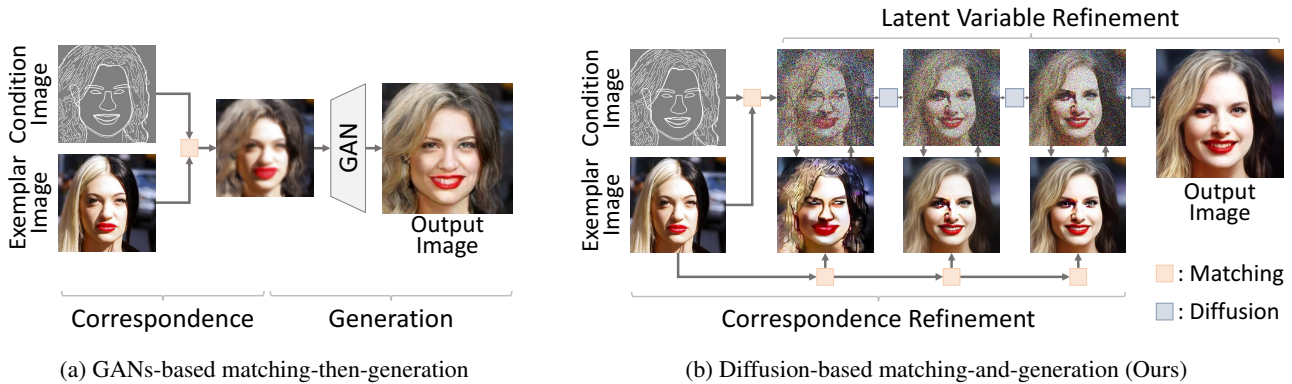


Figure 1: Motivation. (a) existing works (Liao et al. 2017; Zhang et al. 2020; Zhan et al. 2021a,b; Zhou et al. 2021; Zhan et al. 2022) and (b) our MIDMs include the interleaved process of the matching and generation, which can refine correspondence and embedded feature simultaneously.

itation of previous methods (Zhang et al. 2020; Zhou et al. 2021; Zhan et al. 2022, 2021a) while transferring the detail of exemplars faithfully and preserving the structure of condition images.

Experiments demonstrate that our MIDMs achieve competitive performance on CelebA-HQ (Liu et al. 2015) and DeepFashion (Liu et al. 2016). In particular, user study and qualitative comparison results demonstrate that our method can provide a better realistic appearance while capturing the exemplar’s details. An extensive ablation study shows the effectiveness of each component in MIDMs.

Related Work

Exemplar-based Image Translation. There have been a number of works (Bansal, Sheikh, and Ramanan 2019; Wang et al. 2019; Qi et al. 2018; Huang et al. 2018) for exemplar-based image translation. Early works (Huang et al. 2018) focused on bringing global styles, but recent works (Liao et al. 2017; Zhang et al. 2020; Zhan et al. 2021a,b; Zhou et al. 2021; Zhan et al. 2022) have emerged to reference local styles by combining matching networks. While deep image analogy (DIA) (Liao et al. 2017) proposed establishing dense correspondence, CoCosNet (Zhang et al. 2020) suggested that building dense correspondence to cross-domain inputs makes the generated image preserve the given exemplar’s fine details. Followed by this work, CoCosNet v2 (Zhou et al. 2021) integrates PatchMatch (Barnes et al. 2009). Although UNITE (Zhan et al. 2021a) suggested unbalanced optimal transport (Villani 2009) for feature matching to solve the many-to-one alignment problems, establishing feature alignment in cross-domain often fails because of domain gaps. To solve this problem, MCL-Net (Zhan et al. 2022) introduced marginal contrastive loss (Van den Oord, Li, and Vinyals 2018) to explicitly learn the domain-invariant features.

Denoising Diffusion Probabilistic Models. Diffusion models generate a realistic image through the reverse of the noising process. With compelling generation results of many recent studies (Ho, Jain, and Abbeel 2020; Dhariwal and Nichol 2021; Nichol et al. 2021; Rombach et al.

2021; Ramesh et al. 2022), diffusion models have emerged as a competitor to GAN-based generative models. Recently, DDIM (Song, Meng, and Ermon 2020) converted the sampling process to a non-Markovian process, enabling fast and deterministic sampling. Latent diffusion models (LDM) (Rombach et al. 2021) trained the diffusion model in a latent space by adopting a frozen pretrained encoder-decoder structure, which reduces computational complexity.

Meanwhile, conditioning these diffusion models have been studied to make the controllable generation. In SDEdit (Meng et al. 2021), proper amounts of noise were added to a drawing and denoised to recover the realistic image by the reverse process. DiffusionCLIP (Kim and Ye 2021) encodes the input image by the forward process of DDIM and finetunes the diffusion network with text-guided CLIP (Radford et al. 2021) loss. However, there was no study to consider the connection between dense correspondence and image generation based on the diffusion models for exemplar-based image translation, which is the topic of this paper.

Correspondence Learning. Establishing visual correspondences enables building a dense correlation between visually or semantically similar images. Thanks to the rapid advance of convolutional neural networks (CNNs), many works (Long, Zhang, and Darrell 2014; Rocco, Arandjelović, and Sivic 2017; Kim et al. 2017, 2018; Cho et al. 2021; Cho, Hong, and Kim 2022) have shown promising results to estimate semantic correspondence. Incorporating the correspondence model into the diffusion model is the topic of this paper.

Preliminaries

Diffusion Models. Diffusion models enable generating a realistic image from a normal distribution by reversing a gradual noising process (Sohl-Dickstein et al. 2015; Ho, Jain, and Abbeel 2020). Forward process, $q(\cdot)$, is a Markov chain that gradually converts to Gaussian distribution from the data $x_0 \sim q(x_0)$. One step of forward process is defined as $q(x_t|x_{t-1}) := \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t\mathbf{I})$, where β_t is a pre-defined variance schedule in T steps. The forward pro-




Figure 2: Overall architecture of MIDMs. For condition image $I_{\mathcal{X}}$ and exemplar image $I_{\mathcal{Y}}$, we first compute initial matching and obtain the initial warped feature $\mathcal{R}_{\mathcal{X} \leftarrow \mathcal{Y}}$. Then we iteratively compute the diffusion and in-domain alignment with warped feature $r_{\mathcal{Y}}^n$ and reference \mathcal{Y} to finally achieve $r_{\mathcal{Y}}^0$ that is used to achieve $I_{\mathcal{X} \leftarrow \mathcal{Y}}$.

cess can sample x_t at an arbitrary timestamp t in a closed form:

$$\begin{aligned} x_t &= \sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon, \\ \alpha_t &:= \prod_{s=1}^t (1 - \beta_s), \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}). \end{aligned} \quad (1)$$

In addition, the reverse process is defined as $p_{\theta}(x_{t-1}|x_t) := \mathcal{N}(x_{t-1}; \mu_{\theta}(x_t, t), \sigma_{\theta}(x_t, t)\mathbf{I})$ that can be parameterized using deep neural network. DDPMs (Ho, Jain, and Abbeel 2020) found that using noise approximation model $\epsilon_{\theta}(x_t, t)$ worked best instead of using $\mu_{\theta}(x_t, t)$ to procedurally transform the prior noise into data. Therefore, sampling of diffusion models is performed such that

$$x_{t-1} = \frac{1}{\sqrt{1 - \beta_t}} \left(x_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}} \epsilon_{\theta}(x_t, t) \right) + \sigma_t \epsilon. \quad (2)$$

Latent Diffusion Models. Recently, Latent Diffusion Models (LDM) (Rombach et al. 2021) reduces computation cost by learning diffusion model in a latent space. It adopts pretrained encoder \mathcal{E} to embed an image to latent space and pretrained decoder \mathcal{D} to reconstruct the image. In LDM, instead of x itself, $z = \mathcal{E}(x)$ is used to define a diffusion process. Since DDIM (Song, Meng, and Ermon 2020) uses an Euler discretization of some neural ODE (Chen et al. 2018), enabling fast and deterministic sampling, LDM also adopted the DDIM sampling process. Intuitively, the DDIM sampler predicts z_0 directly from z_t and then generates z_{t-1} through a reverse conditional distribution. In specific, $f_{\theta}(z_t, t)$ is a prediction of z_0 given z_t and t :

$$f_{\theta}(z_t, t) := \frac{z_t - \sqrt{1 - \alpha_t} \epsilon_{\theta}(z_t, t)}{\sqrt{\alpha_t}}. \quad (3)$$

The deterministic sampling process of DDIM in LDM is then as follows:

$$z_{t-1} = \sqrt{\alpha_{t-1}} f_{\theta}(z_t, t) + \sqrt{1 - \alpha_{t-1}} \epsilon_{\theta}(z_t, t). \quad (4)$$

After the diffusion process, an image is recovered such that $x = \mathcal{D}(z)$.

On the other hand, numerous works (Saharia et al. 2021; Rombach et al. 2021) proposed a way to condition to the diffusion models. In specific, LDM proposes conditional generation by augmenting diffusion U-Net (Ronneberger, Fischer, and Brox 2015). But these conditioning techniques

cannot be directly applied to exemplar-based image translation tasks, which is the topic of this paper.

Methodology

Problem Statement

Let us denote a condition image and exemplar image as $I_{\mathcal{X}}$ and $I_{\mathcal{Y}}$, e.g., a segmentation map and a real image, respectively. Our objective is to generate an image $I_{\mathcal{X} \leftarrow \mathcal{Y}}$ that follows the content of $I_{\mathcal{X}}$ and the style of $I_{\mathcal{Y}}$, which is called an exemplar-based image translation task.

Conventional works (Zhang et al. 2020; Zhan et al. 2021a, 2022) that solved this task typically followed two steps: cross-domain matching step between input images $I_{\mathcal{X}}$ and $I_{\mathcal{Y}}$ and image generation step from the warping hypothesis. Specifically, they first extract domain-invariant features $S_{\mathcal{X}}$ and $S_{\mathcal{Y}}$ from $I_{\mathcal{X}}$ and $I_{\mathcal{Y}}$, respectively, match them, and estimate an intermediate warp $R_{\mathcal{X} \leftarrow \mathcal{Y}}$ through the matches. An image generator, especially based on GANs (Goodfellow et al. 2014), then generates an output image $I_{\mathcal{X} \leftarrow \mathcal{Y}}$ from $R_{\mathcal{X} \leftarrow \mathcal{Y}}$. However, directly estimating cross-domain correspondence (e.g., sketch-photo) is much more complicated and erroneous than in-domain correspondence. Thus they showed limited performance (Zhan et al. 2021b) depending on the quality of intermediate warped result $R_{\mathcal{X} \leftarrow \mathcal{Y}}$. In addition, they inherit the limitations of GANs, such as less diversity or mode drop in the output distribution (Metz et al. 2016).

Matching Interleaved Diffusion Models (MIDMs)

To alleviate the aforementioned limitations of existing works (Zhang et al. 2020; Zhan et al. 2021a, 2022), as illustrated in Fig. 2, we introduce matching interleaved diffusion models (MIDMs) that interleave cross-domain matching and diffusion steps to modify the diffusion trajectory towards more faithful image translation, i.e., in a warping-and-generation framework. Our framework consists of cross-domain matching and diffusion model-based generation modules, and they are formulated in an iterative manner. In the following, we first explain cross-domain matching and warping, diffusion model-based generation, and their integration in an iterative fashion.




Figure 3: Examples of iterative matching-and-generation process. (a) exemplar image, (b) condition image, (c)-(e) intermediate results of iterative process, which are refined gradually ($n = 5, 4, 2$), and (f) final synthesis result ($n = 0$).

Similarly to LDM (Rombach et al. 2021), we first define cross-domain correspondence and diffusion process in the intermediate latent space from pretrained frozen encoder-decoder, consisting of encoder \mathcal{E} and decoder \mathcal{D} , so as to reduce the computation burden while preserving the image generation quality. In specific, given the condition image $I_{\mathcal{X}}$ and exemplar image $I_{\mathcal{Y}}$, we extract the embedding features $D_{\mathcal{X}}$ and $D_{\mathcal{Y}}$, respectively, through the pretrained encoder (Esser, Rombach, and Ommer 2021) such that $D_{\mathcal{X}} = \mathcal{E}(I_{\mathcal{X}})$ and $D_{\mathcal{Y}} = \mathcal{E}(I_{\mathcal{Y}})$. We abbreviate these as condition and exemplar respectively for the following explanations.

Cross-Domain Correspondence and Warping. For the cross-domain correspondence, our framework reduces a domain discrepancy by introducing two additional encoders, $\mathcal{F}_{\mathcal{X}}$ and $\mathcal{F}_{\mathcal{Y}}$ for condition and exemplar with separated parameters, respectively, to extract common features such that $S_{\mathcal{X}} = \mathcal{F}_{\mathcal{X}}(D_{\mathcal{X}})$ and $S_{\mathcal{Y}} = \mathcal{F}_{\mathcal{Y}}(D_{\mathcal{Y}})$. To estimate the warping hypothesis, we compute a correlation map $C_{\mathcal{X} \leftarrow \mathcal{Y}}$ defined such that

$$C_{\mathcal{X} \leftarrow \mathcal{Y}}(u, v) = \frac{S_{\mathcal{X}}(u)}{\|S_{\mathcal{X}}(u)\|} \cdot \frac{S_{\mathcal{Y}}(v)}{\|S_{\mathcal{Y}}(v)\|}, \quad (5)$$

where u and v index the condition and exemplar features, respectively.

By taking the softmax operation, we can softly warp the exemplar $D_{\mathcal{Y}}$ according to $C_{\mathcal{X} \leftarrow \mathcal{Y}}$:

$$R_{\mathcal{X} \leftarrow \mathcal{Y}}(u) = \sum_v \text{softmax}_v(C_{\mathcal{X} \leftarrow \mathcal{Y}}(u, v)/\tau) D_{\mathcal{Y}}(v), \quad (6)$$

where τ is a temperature, controlling the sharpness of softmax operation.

Latent Variable Refinement Using Diffusion Prior. In this section, we utilize the diffusion process to refine the warped feature. Intuitively, given an initially-warped one, we add an appropriate amount of noise according to the standard forward process of DDPMs (Ho, Jain, and Abbeel 2020) to soften away the unwanted artifacts and distortions which may stem from unreliable correspondences, while preserving the structural information of the warped feature. Specifically, in the diffusion process, we feed $R_{\mathcal{X} \leftarrow \mathcal{Y}}$ to forward the process of DDPMs (Ho, Jain, and Abbeel 2020) to some extent and get the noisy latent variable $r_{\mathcal{Y}}^N$ with proper N . We then iteratively denoise this, following an accelerated

generation process in (Song, Meng, and Ermon 2020):

$$r_{\mathcal{Y}}^n = \begin{cases} \tilde{r}_{\mathcal{Y}}^1 & (n = 0) \\ \sqrt{\alpha_{\tau_N}} R_{\mathcal{X} \leftarrow \mathcal{Y}} + \sqrt{1 - \alpha_{\tau_N}} \epsilon & (n = N) \\ \sqrt{\alpha_{\tau_n}} \tilde{r}_{\mathcal{Y}}^{n+1} + \sqrt{1 - \alpha_{\tau_n}} \epsilon_{\theta}(r_{\mathcal{Y}}^{n+1}, \tau_{n+1}) & (o.w.) \end{cases} \quad (7)$$

where $\tilde{r}_{\mathcal{Y}}^n = f_{\theta}(r_{\mathcal{Y}}^n, \tau_n)$, $\epsilon \sim \mathcal{N}(0, \mathbf{I})$, and *o.w.* means *otherwise*. $\{\tau_n\}$ is a subsequence of time steps in the reverse process, i.e., the number of entire steps in the reverse process is reduced to T , which is the length of $\{\tau_n\}$. $N \in (0, T)$ is an intermediate step to initiate the reverse process. By forwarding diffusion U-net (Rombach et al. 2021) and matching module iteratively, we get the refined latent variable $r_{\mathcal{Y}}^0$.

Interleaving Correspondence and Reverse Process. In this section, we explain how cross-domain correspondence is interleaved with denoising steps in an iterative manner. The intuition behind this is that matching the warped image and exemplar image is more robustly established than the matching between initial content and exemplar images as done in existing methods (Zhang et al. 2020; Zhan et al. 2021a,b; Zhou et al. 2021). Specifically, we first feed the initially-warped exemplar $R_{\mathcal{X} \leftarrow \mathcal{Y}}$ to a noising process to get $r_{\mathcal{Y}}^N$. We then feed it to *one* step of sampling process to get a fully denoised prediction $\tilde{r}_{\mathcal{Y}}^N$. Note that thanks to non-Markovian property of DDIM in Eq. 3, we can directly get a fully denoised prediction $\tilde{r}_{\mathcal{Y}}^N$. In our framework interleaving correspondence and diffusion process, we intercept this, generate a better warped one, and then return to the denoising trajectory using the posterior distribution in Eq. 4.

In this framework, to achieve better correspondence at each step, we compute the correlation between $S_{\mathcal{Y}}$ and $\tilde{r}_{\mathcal{Y}}^n$. To this end, we extract a feature defined such that

$$S_{\mathcal{Y}}^{\text{iter}} = \mathcal{F}_{\mathcal{Y}}^{\text{iter}}(\tilde{r}_{\mathcal{Y}}^n, D_{\mathcal{X}}), \quad (8)$$

where $\mathcal{F}_{\mathcal{Y}}^{\text{iter}}(\cdot)$ is a feature extractor designed for iteration, which receives refined warped exemplar $\tilde{r}_{\mathcal{Y}}^n$ and condition $D_{\mathcal{X}}$ and mixes using a spatially-adaptive normalization (Park et al. 2019). In fact, one can feed the $\tilde{r}_{\mathcal{Y}}^n$ to $\mathcal{F}_{\mathcal{Y}}$ instead of $\mathcal{F}_{\mathcal{Y}}^{\text{iter}}$ since $\tilde{r}_{\mathcal{Y}}^n$ is also from a real distribution. Nevertheless, as shown in (Zhu et al. 2020a), we observe that injecting the condition $D_{\mathcal{X}}$ into the feature extractor can help to align the features and build more correct correspondences. We then compute a correlation map $C_{\mathcal{X} \leftarrow \mathcal{Y}}^{\text{iter}}$ with $S_{\mathcal{Y}}$ and $S_{\mathcal{Y}}^{\text{iter}}$ and extract the $R_{\tilde{r}_{\mathcal{Y}}^n \leftarrow \mathcal{Y}}$. By returning $R_{\tilde{r}_{\mathcal{Y}}^n \leftarrow \mathcal{Y}}$ to

the denoising trajectory according to Eq. 4, we can obtain r_y^{n-1} . By iterating the above process, we finally obtain r_y^0 . To summary, for $1 \leq n < N$, we change \tilde{r}_y^{n+1} in Eq. 7 to $R_{\tilde{r}_y^{n+1} \leftarrow y}$ as follows:

$$r_y^n = \sqrt{\alpha_{\tau_n}} R_{\tilde{r}_y^{n+1} \leftarrow y} + \sqrt{1 - \alpha_{\tau_n}} \epsilon_{\theta}(r_y^{n+1}, \tau_{n+1}). \quad (9)$$

Confidence-Aware Matching. There is a trade-off between bringing the details of exemplar faithfully and generating an image that matches the condition image, e.g., in the case of the condition image having earrings that do not exist in exemplar image (Zhang et al. 2020). To address this problem, we additionally propose a confidence-based masking technique. Specifically, we utilize a cycle-consistency (Jiang et al. 2021) as the matching confidence at each warping step. We define the confidence mask such that

$$M_{\tilde{r}_y^n \leftarrow y}(u) = \mathbf{1}(\|u - \psi_{y \leftarrow \tilde{r}_y^n}(\psi_{\tilde{r}_y^n \leftarrow y}(u))\|_2^2 < \gamma) \quad (10)$$

where ψ is a warping function (Jiang et al. 2021) and γ is a threshold constant. Using this confidence mask $M_{\tilde{r}_y^n \leftarrow y}$, we only rewrap the confident region and the rest region skips the rewrapping process in Eq. 9 as

$$r_y^n = \sqrt{\alpha_{\tau_n}} (M_{\tilde{r}_y^n \leftarrow y} \odot R_{\tilde{r}_y^{n+1} \leftarrow y} + (1 - M_{\tilde{r}_y^n \leftarrow y}) \odot \tilde{r}_y^{n+1}) + \sqrt{1 - \alpha_{\tau_n}} \epsilon_{\theta}(r_y^{n+1}, \tau_{n+1}), \quad (11)$$

for $1 \leq n < N$. With this technique, the regions with low matching confidence intend to follow the reverse process of the general diffusion model. Intuitively, it allows selective control of the generative power depending on the matching confidence of the regions, which alleviates the aforementioned problem.

Image Reconstruction. Finally, we get the translated images by returning the latent variables to image space such that $I_{\mathcal{X} \leftarrow y} = \mathcal{D}(r_y^0)$. We illustrate the whole process described above in Fig. 3.

Loss Functions

Our model incorporates several losses to accomplish photo-realistic image translation. Note that we fine-tune the diffusion model with our loss functions.

Losses for Cross-Domain Correspondence. We use a pseudo-ground-truth image of a condition input image $I_{\mathcal{X}}$ as $I'_{\mathcal{X}}$. We need to ensure that the extracted common features $\mathcal{S}_{\mathcal{X}}$ and $\mathcal{S}'_{\mathcal{X}}$ are in the same domain.

$$\mathcal{L}_{\text{dom}} = \|\mathcal{S}'_{\mathcal{X}} - \mathcal{S}_{\mathcal{X}}\|_1. \quad (12)$$

In addition, the warped features should be cycle-consistent, which means that the exemplar needs to be returnable from the warped features. Because of our interleaved warping and generation process, we can acquire the cyclic-warped features at every n -th step:

$$\mathcal{L}_{\text{cycle}} = \sum_n \|R_{y \leftarrow \tilde{r}_y^{n+1} \leftarrow y} - D_y\|_1, \quad (13)$$

where $R_{y \leftarrow \tilde{r}_y^{n+1} \leftarrow y}$ is the cyclic-warped reference feature at n -step.



Figure 4: Qualitative results for edge-to-face on CelebA-HQ (Liu et al. 2015): (from top to bottom) exemplars, condition and results by CoCosNet (Zhang et al. 2020) and our MIDMs.

Finally, when we warp the ground-truth feature $\mathcal{D}'_{\mathcal{X}}$ with the correlation $\mathcal{C}_{\mathcal{X} \leftarrow \mathcal{X}'}$, we can obtain $R_{\mathcal{X} \leftarrow \mathcal{X}'}$, and this is consistent in terms of semantics, with the original ground-truth feature $\mathcal{D}'_{\mathcal{X}}$, building a source-condition loss \mathcal{L}_{src} as specified below:

$$\mathcal{L}_{\text{src}} = \|\phi_l(I_{\mathcal{X} \leftarrow \mathcal{X}'} - \phi_l(I'_{\mathcal{X}}))\|_1, \quad (14)$$

where ϕ_l is a l -th activation layer of pretrained VGG-19 model (Simonyan and Zisserman 2015).

Losses for Image-to-image Translation. We use a perceptual loss (Johnson, Alahi, and Fei-Fei 2016) to maximize the semantic similarity since the semantic of the produced image should be consistent with the conditional input $I_{\mathcal{X}}$ or the ground truth $I'_{\mathcal{X}}$, denoted as follows:

$$\mathcal{L}_{\text{perc}} = \|\phi_l(I_{\mathcal{X} \leftarrow y} - \phi_l(I'_{\mathcal{X}}))\|_1. \quad (15)$$

Besides, we encourage the generated image $I_{\mathcal{X} \leftarrow y}$ to take the style consistency with the semantically corresponding patches from the exemplar $I'_{\mathcal{X}}$. Thus, we choose the contextual loss (Mechrez, Talmi, and Zelnik-Manor 2018) as a style loss, expressed in the form of:

$$\mathcal{L}_{\text{style}} = -\log \left(\sum_l \mu_i \text{CX}_{ij}(\phi_l(I_{\mathcal{X} \leftarrow y}), \phi_l(I_y)) \right) \quad (16)$$

where CX_{ij} is a contextual similarity function between images (Mechrez, Talmi, and Zelnik-Manor 2018).

Loss for Diffusion. We fine-tune a pretrained diffusion model (Rombach et al. 2021). The diffusion objectives are defined as:

$$\mathcal{L}_{\text{diff}} = \sum_n \|\epsilon_{\theta}(r_y^{n+1}, \tau_{n+1}) - \epsilon\|_2, \quad (17)$$

where ϵ is random noise used in the forward process of the diffusion (Ho, Jain, and Abbeel 2020).

Experiments

Experimental Settings

Datasets. Following the previous literature (Zhang et al. 2020; Zhan et al. 2021b,a), we conduct experiments over

Methods	DeepFashion (Liu et al. 2016)				CelebA-HQ (Liu et al. 2015)			
	w/ sources			w/ exemplars	w/ sources			w/ exemplars
	FID ↓	SWD ↓	LPIPS ↑	FID ↓	FID ↓	SWD ↓	LPIPS ↑	FID ↓
Pix2pixHD (Wang et al. 2018)	25.20	16.40	-	-	42.70	33.30	-	-
SPADE (Park et al. 2019)	36.20	27.80	0.231	-	31.50	26.90	0.187	-
SelectionGAN (Tang et al. 2019)	38.31	28.21	0.223	-	34.67	27.34	0.191	-
SMIS (Zhu et al. 2020b)	22.23	23.73	0.240	-	23.71	22.23	0.201	-
SEAN (Zhu et al. 2020a)	16.28	17.52	0.251	-	18.88	19.94	0.203	-
UNITE (Zhan et al. 2021a)	13.08	16.65	0.278	-	13.15	14.91	0.213	-
CoCosNet (Zhang et al. 2020)	14.40	17.20	0.272	11.12	14.30	15.30	0.208	11.01
CoCosNet v2 (Zhou et al. 2021)	12.81	16.53	0.283	-	12.85	14.62	0.218	-
MCL-Net (Zhan et al. 2022)	12.89	16.24	0.286	-	12.52	14.21	0.216	-
MIDMs (Ours)	10.89	10.10	0.279	8.54	15.67	12.34	0.224	10.67

Table 1: Quantitative evaluation on DeepFashion and CelebA-HQ datasets. The comparisons are performed with three widely used evaluation metrics FID (Heusel et al. 2017), SWD (Karras et al. 2017) and LPIPS (Zhang et al. 2018). Some metrics for method that cannot be experimented because the codes or trained weights are not available are left blank.




Figure 5: Qualitative results for keypoints-to-photos on DeepFashion (Liu et al. 2016). (from top to bottom) exemplars, condition and results by CoCosNet (Zhang et al. 2020) and our MIDMs.

the CelebA-HQ (Liu et al. 2015), and DeepFashion (Liu et al. 2016) datasets. CelebA-HQ (Liu et al. 2015) dataset provides 30,000 images of high-resolution human faces at 1024×1024 resolution, and we construct the edge maps using Canny edge detector (Canny 1986) for conditional input. DeepFashion (Liu et al. 2016) dataset consists of 52,712 full-length person images in fashion cloths with the keypoints annotations obtained by OpenPose (Cao et al. 2021). Also, we use LSUN-Churches (Yu et al. 2015) to conduct the experiments of segmentation maps-to-photos. In the case of LSUN-Churches dataset, we generate segmentation maps using Swin-S (Liu et al. 2021) trained on ADE20k (Zhou et al. 2017).

Implementation Details. We use AdamW optimizer (Loshchilov and Hutter 2017) for the learning rate of $3e-6$ for the correspondence network, and $1.5e-7$ for the backbone network of the diffusion model. We use multi-step learning rate decay with $\gamma = 0.3$. We conduct our all experiments on RTX 3090 GPU, and we provide more implementation details and pseudo code in the Appendix.




Figure 6: Qualitative results for segmentation maps-to-photos on LSUN-Churches (Yu et al. 2015).

Methods	Style relevance		Semantic consistency
	Color	Texture	
Pix2PixHD (Wang et al. 2018)	-	-	0.914
SPADE (Park et al. 2019)	0.955	0.927	0.922
MUNIT (Huang et al. 2018)	0.939	0.884	0.848
EGSC-IT (Ma et al. 2018)	0.965	0.942	0.915
CoCosNet (Zhang et al. 2020)	0.977	0.958	0.949
MIDMs (Ours)	0.982	0.962	0.915

Table 2: Quantitative evaluation of style relevance and semantic consistency on CelebA-HQ (Liu et al. 2015).

Evaluation Metrics. To evaluate the translation results comprehensively, we report Fréchet Inception Score (FID) (Heusel et al. 2017) and Sliced Wasserstein distance (SWD) to evaluate the image perceptual quality, (Karras et al. 2017), and Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al. 2018) scores to evaluate the diversity of translated images. Furthermore, we employ the style relevance and semantic consistency metrics (Zhang et al. 2020) using a pretrained VGG model (Simonyan and Zisserman 2015), which measures the cosine similarity between features of translated results and exemplar inputs. Specifically, the low-level features (*i.e.*, outputs of pretrained VGG network at `relu1_2` and `relu2_2` layers) are used to calculate color and style relevance, and high-level




Figure 7: User studies on CelebA-HQ (Liu et al. 2015) and DeepFashion (Liu et al. 2016).

features (*i.e.*, outputs of `relu3_2`, `relu4_2` and `relu5_2` layers) are used to compute the semantic consistency score.

Qualitative Evaluation

Fig. 4 and Fig. 5 demonstrate qualitative results with respect to different condition styles compared to CoCosNet (Zhang et al. 2020). As can be seen therein, our method translates the detailed style of exemplar well in both datasets, preserving the structures of condition images. We also show diverse results on LSUN-Churches (Yu et al. 2015) in Fig. 6.

Quantitative Evaluation

Table 1 shows quantitative comparison with other exemplar-based image translation methods. MIDMs outperform with large gaps in terms of SWD in both datasets. Also in other metrics, our method demonstrates superior or competitive performance. The semantic consistency and style consistency performance evaluations are summarized in Table 2. The proposed method achieves the best style relevance scores including both color and texture. We additionally evaluate the FID score compared with not only the distribution of source images as prior works (Zhang et al. 2020; Van den Oord, Li, and Vinyals 2018) did, but also the distribution of exemplar images. In terms of FID compared with the distribution of exemplar images, MIDMs show superior results on all datasets we experiment with, which can be seen that our method translates the style of exemplar better.

User Study

Finally, we conduct a user study to compare the subjective quality of the translated results in Fig. 7. From 61 participants, we ask to rank all the methods in terms of style relevance and quality. In all studies, we outperform the CoCosNet (Zhang et al. 2020) and UNITE (Zhan et al. 2021a).

Ablation Study

We conduct ablation studies to demonstrate our model can find better correspondence generating more realistic images. Also, more ablation studies can be found in the Appendix.

Network Designs. From our best model, we validate our contribution by taking out the components of our model one by one in Table 3. We observe a consistent decrease in performance when each component is removed. We can find that confidence masking is effective for our model. Replacing the recurrent matching process with one-time matching

Models	FID↓	SWD↓
Ours	15.67	12.34
(-) Confidence Masking	19.21	16.01
(-) Recurrent Matching	24.76	23.71
(-) Diffusion U-Net	128.70	34.59

Table 3: Ablation study on the variants of components. The baseline is our best model, and we validate the performance on CelebA-HQ (Liu et al. 2015) by removing the elements one by one.

Noise	FID↓
20%	23.67
25%	15.67
30%	16.01
35%	19.20

Table 4: Performance with respect to the noise levels at sampling. We evaluate the performance on CelebA-HQ (Liu et al. 2015).

Loss	FID↓
Ours	15.67
w/o $\mathcal{L}_{\text{cycle}}$	16.18
w/o $\mathcal{L}_{\text{style}}$	19.23
w/o $\mathcal{L}_{\text{perc}}$	16.51
w/o \mathcal{L}_{dom}	16.68
w/o \mathcal{L}_{src}	72.25

Table 5: Ablation study on each loss function. We evaluate the performance on CelebA-HQ (Liu et al. 2015).

degrades the image quality significantly, which proves the superiority of our approach compared to the matching-then-generation framework.

Evaluations on Different Noise Levels. We also evaluate the FID score of our model for the different noise labels, and the results are shown in Table 4. We observe that the proposed method with the 25% noise level shows the best performance.

Loss Functions. We conduct an ablation study to confirm the performance contribution of each loss function, by removing the loss term from our overall loss functions, and the result is shown in Table 5:

Conclusion

In this paper, we presented MIDMs that interleave cross-domain matching and diffusion steps in the latent space by iteratively feeding the intermediate warp into the noising process and denoising it to generate a translated image. To the best of our knowledge, it is the first attempt to use the diffusion models as a competitor to GANs-based methods in exemplar-based image translation. Thanks to the joint synergy of the proposed modules, the style of exemplar were reliably translated to the condition input. Experimental results show the superiority of our MIDMs for exemplar-based image translation as well as a general image translation task.

Acknowledgments

This research was supported by the MSIT, Korea (IITP-2023-2020-0-01819, ICT Creative Consilience program, No.2021-0-00155, Context and Activity Analysis-based Solution for Safe Childcare), and National Research Foundation of Korea (NRF-2021R1C1C1006897).

References

- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein generative adversarial networks. In *ICML*, 214–223. PMLR.
- Bansal, A.; Sheikh, Y.; and Ramanan, D. 2019. Shapes and Context: In-The-Wild Image Synthesis & Manipulation. In *CVPR*, 2312–2321.
- Barnes, C.; Shechtman, E.; Finkelstein, A.; and Goldman, D. B. 2009. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3): 24.
- Brock, A.; Donahue, J.; and Simonyan, K. 2018. Large scale GAN training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*.
- Canny, J. 1986. A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6): 679–698.
- Cao, Z.; Hidalgo, G.; Simon, T.; Wei, S.-E.; and Sheikh, Y. 2021. OpenPose: Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43: 172–186.
- Chen, Q.; and Koltun, V. 2017. Photographic image synthesis with cascaded refinement networks. In *ICCV*, 1511–1520.
- Chen, T. Q.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. K. 2018. Neural Ordinary Differential Equations. *ArXiv*, abs/1806.07366.
- Cho, S.; Hong, S.; Jeon, S.; Lee, Y.; Sohn, K.; and Kim, S. 2021. CATs: Cost Aggregation Transformers for Visual Correspondence. In *NeurIPS*.
- Cho, S.; Hong, S.; and Kim, S. 2022. CATs++: Boosting Cost Aggregation with Convolutions and Transformers. *arXiv preprint arXiv:2202.06817*.
- Choi, J.; Kim, S.; Jeong, Y.; Gwon, Y.; and Yoon, S. 2021. ILVR: Conditioning Method for Denoising Diffusion Probabilistic Models. In *ICCV*, 14367–14376.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion Models Beat GANs on Image Synthesis. In *NeurIPS*, volume 34.
- Esser, P.; Rombach, R.; and Ommer, B. 2021. Taming Transformers for High-Resolution Image Synthesis. In *CVPR*, 12868–12878.
- Goodfellow, I. 2016. Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *NeurIPS*, volume 27.
- Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; and Courville, A. C. 2017. Improved training of wasserstein gans. In *NeurIPS*, volume 30.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, volume 30.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In *NeurIPS*, volume 33, 6840–6851.
- Huang, X.; Liu, M.-Y.; Belongie, S.; and Kautz, J. 2018. Multimodal Unsupervised Image-to-image Translation. In *ECCV*, 172–189.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image Translation with Conditional Adversarial Networks. In *CVPR*, 1125–1134.
- Jiang, W.; Trulls, E.; Hosang, J. H.; Tagliasacchi, A.; and Yi, K. M. 2021. COTR: Correspondence Transformer for Matching Across Images. In *ICCV*, 6187–6197.
- Johnson, J.; Alahi, A.; and Fei-Fei, L. 2016. Perceptual Losses for Real-Time Style Transfer and Super-Resolution. In *ECCV*.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2017. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.
- Kim, G.; and Ye, J. C. 2021. Diffusionclip: Text-guided Image Manipulation using Diffusion models. *arXiv preprint arXiv:2110.02711*.
- Kim, S.; Lin, S.; Jeon, S.; Min, D.; and Sohn, K. 2018. Recurrent Transformer Networks for Semantic Correspondence. In *NeurIPS*.
- Kim, S.; Min, D.; Ham, B.; Jeon, S.; Lin, S.; and Sohn, K. 2017. FCSS: Fully Convolutional Self-Similarity for Dense Semantic Correspondence. In *CVPR*, 616–625.
- Liao, J.; Yao, Y.; Yuan, L.; Hua, G.; and Kang, S. B. 2017. Visual attribute transfer through deep image analogy. *ACM Transactions on Graphics (TOG)*, 36: 1 – 15.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10012–10022.
- Liu, Z.; Luo, P.; Qiu, S.; Wang, X.; and Tang, X. 2016. Deepfashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations. In *CVPR*, 1096–1104.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep Learning Face Attributes in the Wild. In *ICCV*, 3730–3738.
- Long, J. L.; Zhang, N.; and Darrell, T. 2014. Do convnets learn correspondence? In *NeurIPS*, volume 27.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Ma, L.; Jia, X.; Georgoulis, S.; Tuytelaars, T.; and Van Gool, L. 2018. Exemplar Guided Unsupervised Image-to-image Translation with Semantic Consistency. *arXiv preprint arXiv:1805.11145*.
- Mehrez, R.; Talmi, I.; and Zelnik-Manor, L. 2018. The contextual loss for image transformation with non-aligned data. In *ECCV*, 768–783.

- Meng, C.; Song, Y.; Song, J.; Wu, J.; Zhu, J.-Y.; and Ermon, S. 2021. Sdedit: Image Synthesis and Editing with Stochastic Differential Equations. *arXiv preprint arXiv:2108.01073*.
- Metz, L.; Poole, B.; Pfau, D.; and Sohl-Dickstein, J. 2016. Unrolled generative adversarial networks. *arXiv preprint arXiv:1611.02163*.
- Miyato, T.; Kataoka, T.; Koyama, M.; and Yoshida, Y. 2018. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*.
- Nichol, A.; Dhariwal, P.; Ramesh, A.; Shyam, P.; Mishkin, P.; McGrew, B.; Sutskever, I.; and Chen, M. 2021. Glide: Towards Photorealistic Image Generation and Editing with Text-guided Diffusion Models. *arXiv preprint arXiv:2112.10741*.
- Park, T.; Liu, M.-Y.; Wang, T.-C.; and Zhu, J.-Y. 2019. Semantic Image Synthesis with Spatially-adaptive Normalization. In *CVPR*, 2337–2346.
- Qi, X.; Chen, Q.; Jia, J.; and Koltun, V. 2018. Semi-Parametric Image Synthesis. In *CVPR*, 8808–8816.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; and Sutskever, I. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical Text-conditional Image Generation with CLIP Latents. *arXiv preprint arXiv:2204.06125*.
- Rocco, I.; Arandjelović, R.; and Sivic, J. 2017. Convolutional Neural Network Architecture for Geometric Matching. In *CVPR*, 39–48.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2021. High-Resolution Image Synthesis with Latent Diffusion Models. *arXiv preprint arXiv:2112.10752*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *MICCAI*.
- Saharia, C.; Chan, W.; Chang, H.; Lee, C. A.; Ho, J.; Salimans, T.; Fleet, D. J.; and Norouzi, M. 2021. Palette: Image-to-image Diffusion Models. *arXiv preprint arXiv:2111.05826*.
- Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; and Chen, X. 2016. Improved techniques for training gans. In *NeurIPS*, volume 29.
- Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR*, abs/1409.1556.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep Unsupervised Learning Using Nonequilibrium Thermodynamics. In *ICML*, 2256–2265. PMLR.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising Diffusion Implicit Models. *arXiv preprint arXiv:2010.02502*.
- Tang, H.; Xu, D.; Sebe, N.; Wang, Y.; Corso, J. J.; and Yan, Y. 2019. Multi-channel attention selection gan with cascaded semantic guidance for cross-view image translation. In *CVPR*, 2417–2426.
- Van den Oord, A.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv e-prints*, arXiv–1807.
- Villani, C. 2009. *Optimal transport: old and new*, volume 338. Springer.
- Wang, M.; Yang, G.-Y.; Li, R.; Liang, R.-Z.; Zhang, S.-H.; Hall, P. M.; and Hu, S.-M. 2019. Example-guided Style-consistent Image Synthesis from Semantic Labeling. In *CVPR*, 1495–1504.
- Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018. High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. In *CVPR*, 8798–8807.
- Yu, F.; Seff, A.; Zhang, Y.; Song, S.; Funkhouser, T.; and Xiao, J. 2015. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*.
- Zhan, F.; Yu, Y.; Cui, K.; Zhang, G.; Lu, S.; Pan, J.; Zhang, C.; Ma, F.; Xie, X.; and Miao, C. 2021a. Unbalanced Feature Transport for Exemplar-based Image Translation. In *CVPR*, 15028–15038.
- Zhan, F.; Yu, Y.; Wu, R.; Cui, K.; Xiao, A.; Lu, S.; and Shao, L. 2021b. Bi-level Feature Alignment for Versatile Image Translation and Manipulation. *ArXiv*, abs/2107.03021.
- Zhan, F.; Yu, Y.; Wu, R.; Zhang, J.; Lu, S.; and Zhang, C. 2022. Marginal Contrastive Correspondence for Guided Image Generation. *arXiv preprint arXiv:2204.00442*.
- Zhang, P.; Zhang, B.; Chen, D.; Yuan, L.; and Wen, F. 2020. Cross-domain Correspondence Learning for Exemplar-based Image Translation. In *CVPR*, 5143–5153.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 586–595.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene Parsing through ADE20K Dataset. In *CVPR*, 633–641.
- Zhou, X.; Zhang, B.; Zhang, T.; Zhang, P.; Bao, J.; Chen, D.; Zhang, Z.; and Wen, F. 2021. Cocosnet v2: Full-resolution Correspondence Learning for Image Translation. In *CVPR*, 11465–11475.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired Image-to-image Translation using Cycle-Consistent Adversarial Networks. In *ICCV*, 2223–2232.
- Zhu, P.; Abdal, R.; Qin, Y.; and Wonka, P. 2020a. Sean: Image synthesis with semantic region-adaptive normalization. In *CVPR*, 5104–5113.
- Zhu, Z.; Xu, Z.; You, A.; and Bai, X. 2020b. Semantically multi-modal image synthesis. In *CVPR*, 5467–5476.