

LANIT: Language-Driven Image-to-Image Translation for Unlabeled Data

Jihye Park ^{*1}, Sunwoo Kim ^{*1}, Soohyun Kim ^{*1}, Seokju Cho ¹,
 Jaejun Yoo ², Youngjung Uh ³, Seungryong Kim ^{†1}

¹ Korea University, Seoul, Korea ² UNIST, Ulsan, Korea ³ Yonsei University, Seoul, Korea

¹{ghp1112, sw-kim, shkim1211, seokju.cho, seungryong.kim}@korea.ac.kr
²jaejun.yoo@unist.ac.kr ³yj.uh@yonsei.ac.kr

Abstract

Existing techniques for image-to-image translation commonly have suffered from two critical problems: heavy reliance on per-sample domain annotation and/or inability to handle multiple attributes per image. Recent truly-unsupervised methods adopt clustering approaches to easily provide per-sample one-hot domain labels. However, they cannot account for the real-world setting: one sample may have multiple attributes. In addition, the semantics of the clusters are not easily coupled to human understanding. To overcome these, we present LANGUAGE-driven Image-to-image Translation model, dubbed LANIT. We leverage easy-to-obtain candidate attributes given in texts for a dataset: the similarity between images and attributes indicates per-sample domain labels. This formulation naturally enables multi-hot labels so that users can specify the target domain with a set of attributes in language. To account for the case that the initial prompts are inaccurate, we also present prompt learning. We further present domain regularization loss that enforces translated images to be mapped to the corresponding domain. Experiments on several standard benchmarks demonstrate that LANIT achieves comparable or superior performance to existing models. The code is available at github.com/KU-CVLAB/LANIT.

1. Introduction

Unpaired image-to-translation frameworks aim to translate an image from one domain to another [21, 31, 46]. They are typically trained on many images with their respective domain labels. It has been widely studied for a decade in numerous computer vision applications, such as image manipulation [14, 62, 66], multi-modal translation [3, 9, 10, 21, 35], and instance-aware translation [6, 23, 30].

Conventional methods [9, 10, 21] required at least *per-sample* domain supervision based on the strict assumption that each image should be described by a one-hot domain label, which is frequently violated; e.g., a face can have multiple attributes. In addition, annotating domain label for each image may become considerably labor-intensive and error-prone, and subjectivity may lead to inconsistency across different annotators.

To mitigate such a heavy reliance on per-sample domain annotation, some recent methods formulate the problem in a few-shot setting [35, 51] or semi-supervised learning setting [55]. More recently, TUNIT [3] and Kim *et al.* [29] present a framework for jointly learning domain clustering and translation within a unified model, thus enabling a *truly-unsupervised* learning of image-to-image translation. Although these methods ease the burden of per-sample domain annotation and show relatively competitive performance, they still inherit the limitation of using one-hot domain labels to train the image translation model. In addition, these truly-unsupervised learning methods only learn domain clusters that are dominant in the dataset, which sometimes do not reflect semantic meanings. Such limitations make users confused in understanding the semantic meaning of each learned domain cluster.

In this paper, to overcome the problems in per-sample-level and unsupervised learning methods, we propose using the *dataset-level* annotation for the first time as exemplified in Fig. 1. In practice, it could be much easier to list candidate textual domain descriptions that describe images in the dataset. Using the candidate attributes given in texts for a dataset, we specify the target domain given a reference image according to its similarity to the attributes in vision-language embedding [49]. Then, a generator translates a source image to the target domain using the target style code of the reference image. As similarities naturally lead to a multi-hot label, we can aggregate the style codes of multiple domains to represent the style code for a joint domain. In addition, to account for the fact that the initial prompts may

^{*}Equal contribution

[†]Corresponding author

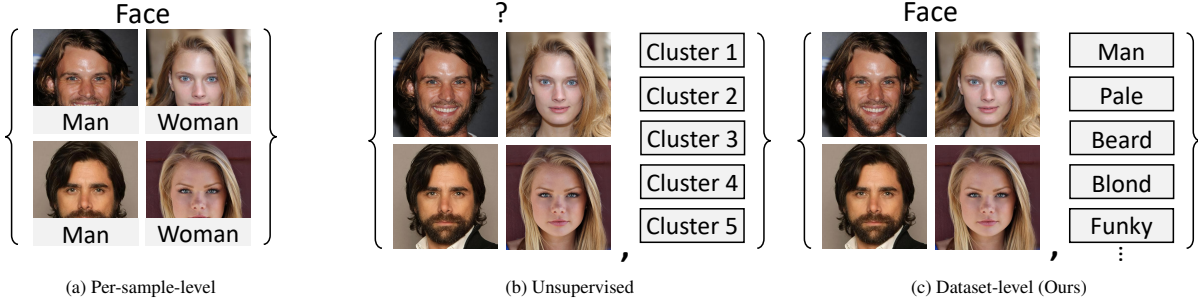


Figure 1. **Levels of supervision.** For unpaired image-to-image translation, (a) conventional methods [10, 21, 35, 55] require at least *per-sample*-level domain supervision, which is often hard to collect. To overcome this, (b) unsupervised learning methods [3, 29] learn image translation model using a dataset itself without any supervision, but it shows limited performance and lacks the semantic understanding of each cluster, limiting its applicability. Unlike them, (c) we present a novel framework that requires a dataset with possible textual domain descriptions (i.e., *dataset-level* annotation), which achieves comparable or even better performance than previous methods.



Figure 2. **Examples of semantic encoding.** The existing unsupervised method [3] allows users to translate one of several clusters. However, the learned clusters lack semantic meaning and sometimes some attributes do not appear in any clusters. Unlike this, our framework can select and train the domains that have explicit semantic meaning, which is more applicable.

be sometimes inaccurate, we also propose prompt learning and encourage the translated images to have corresponding labels with our proposed domain regularization loss.

Our proposed framework just requires some textual domain descriptions for defining domains and it can ease the burden on the per-sample annotating process. Thanks to the proposed dataset-level annotation, users can explicitly define the semantic meaning of the target domain with textual descriptions, as exemplified in Fig. 2. We evaluate our method on several standard benchmarks [7, 26, 35]. Experimental results show that the proposed model trained only with dataset-level supervision is comparable or even outperforms the latest methods, including StarGAN2 [10] trained with per-sample supervision. We also provide intensive ablation studies and user study results.

2. Related Work

Image-to-Image Translation. While early methods for image-to-image translation considered a paired setting [22], most recent state-of-the-arts focus on an unpaired setting [3, 11, 19, 40, 46, 57, 62, 66]. This trend was initiated

by CycleGAN [67], of which many variants were proposed, formulating uni-modal models [34, 60, 61] and multi-modal models [9, 10, 21, 31]. Although these methods overcome the inherent limitation of the paired setting, they still require at least *per-sample* domain annotation, which is often labor intensive or even impossible to achieve in some cases. In addition, this per-sample supervision does not consider that some images may require multi-hot domain labels such as facial attributes, e.g., blond and young.

To mitigate the reliance on per-sample supervision, FUNIT [35] and CoCo-FUNIT [51] proposed a few-shot setting, but still required a few annotations. SEMIT [55] utilizes a pseudo labeling technique. S³GAN [42], Self-conditioned GAN [37], and [4] adopt clustering methods or a classifier. Recently, TUNIT [3] and Kim *et al.* [29] have proposed the *truly-unsupervised* frameworks with clustering approach or prototype-based learning, which do not need any per-sample domain supervision. However, their performance was limited in that it is challenging to understand the semantic meaning of each pseudo domain, which in turn limits their applicability. Note that the aforementioned methods also inherit the problem of the one-hot domain labeling assumption.

Vision-Language Model in Image Manipulation. Many large-scale vision-language models [24, 49] have been popularly adopted in computer vision problems. Especially, CLIP [49] has been employed in image synthesis and manipulation with GANs [1, 38, 47, 56]. For instance, StyleCLIP [47] and HairCLIP [56] manipulate the latent space of StyleGAN [27], through the guidance of CLIP, but they heavily rely on pre-trained StyleGAN. CLIP2StyleGAN extends such an approach in an unsupervised manner that eliminates the need for target style descriptions.

More recently, some works [2, 12, 16, 28, 39, 45, 50] combine a diffusion model [18] with CLIP [49] for text-driven image manipulation with high fidelity. However, they show limited performance in realistically translating in the real image domain. Furthermore, they rely on pre-trained dif-

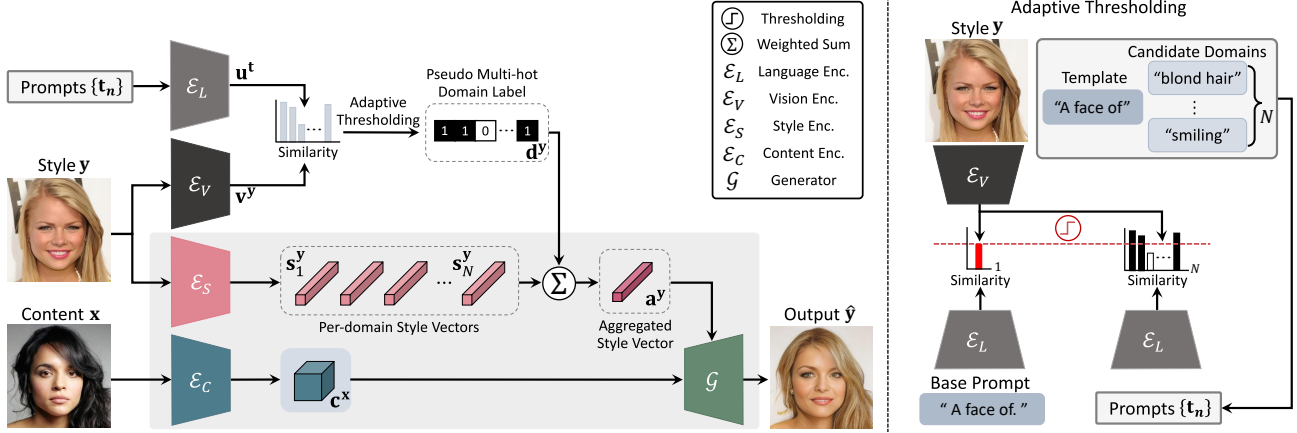


Figure 3. **Network configuration.** Our model consists of content and style encoders (\mathcal{E}_C , \mathcal{E}_S), vision-language encoders (\mathcal{E}_V , \mathcal{E}_L), and generator (\mathcal{G}). We extract content and style vectors from content \mathbf{x} and style \mathbf{y} , respectively. By leveraging vision-language features and the proposed adaptive thresholding technique, we measure the pseudo domain label \mathbf{d}^y of \mathbf{y} . We generate $\hat{\mathbf{y}}$ with content vector \mathbf{c}^x and aggregated style vector \mathbf{a}^y through the generator.

fusion model [18], and they are only available on text-conditioned manipulation. On the other hand, our LANIT is not dependent on pre-trained StyleGAN [27] and also learns style manifold from given image dataset and a few numbers of prompts that designate explicit semantics to be learned.

Prompt Learning. Since the success of pretrained large-scale language models such as GPT-3 [8] in NLP, various methods for optimizing the prompt in a text format have been proposed [33, 36, 48, 52]. Inspired by these works, some works [13, 20, 25, 32, 59, 64, 65] using CLIP [49] also attempted to optimize the input prompt. For instance, CoCoOp [65] showed that using the learned continuous prompt could surpass the manually-designed discrete prompt-based method on zero-shot image classification. However, they require class supervision. Recently, several works [20, 53, 63] propose unsupervised prompt tuning. In this work, we explore combining image translation with prompt learning in an unsupervised way for the first time.

3. Methodology

3.1. Motivation and Problem Formulation

We leverage a list of dataset-level domain descriptions in the text form and their similarity with the unlabeled images in CLIP [49] embedding space to measure a pseudo domain label of style image. By doing so, we aim to alleviate difficulty in the labeling process and give explicit semantic meaning in the learned domain for the multi-hot label setting.

One naive approach obtaining the pseudo domain label is by calculating the similarity between image and text features with CLIP [49]. However, we observe that this leads to inaccurate labeling and noisy translation since the image-

text feature space in CLIP [49] is not perfectly aligned. To mitigate this issue, we introduce two solutions, an adaptive thresholding with the base prompt for selecting confident pseudo labels, and a prompt learning technique with domain regularization loss for boosting confidence in calculating similarity.

3.2. Image Translation with Pseudo Domain Label

Our model, consisting of content encoder \mathcal{E}_C , style encoder \mathcal{E}_S , mapping encoder \mathcal{E}_M , and generator \mathcal{G} similar to StarGAN2 [10], aims to learn an image-to-image translation with content image \mathbf{x} and style image \mathbf{y} (or between domains \mathcal{X} and \mathcal{Y}) to generate a translated image $\hat{\mathbf{y}}$.

In specific, we first extract the content vector and style vector from $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$ and $\mathbf{y} \in \mathbb{R}^{H \times W \times 3}$ with height H and width W such that $\mathbf{c}^x = \mathcal{E}_C(\mathbf{x}) \in \mathbb{R}^{h_c \times w_c \times c}$ with height h_c , width w_c and channel c , and $\mathbf{s}_n^y = \mathcal{E}_{S,n}(\mathbf{y}) \in \mathbb{R}^{1 \times s}$ for n -th domain (or $\tilde{\mathbf{s}} = \mathcal{E}_M(\mathbf{z})$ from a random latent \mathbf{z}), respectively. While existing methods, e.g., StarGAN2 [10], select a single vector from $\{\mathbf{s}_1^y, \dots, \mathbf{s}_N^y\}$ of N domains according to ground-truth one-hot domain label, which is inserted to the generator, we aggregate such style vectors $\{\mathbf{s}_1^y, \dots, \mathbf{s}_N^y\}$ with the pseudo domain label $\mathbf{d}^y = [d_n^y]_{n=1}^N \in \mathbb{R}^{N \times 1}$ defined as a binary vector, which will be discussed in the following, such that

$$\mathbf{a}^y = \frac{1}{M^y} \sum_{n=1}^N \mathbf{s}_n^y d_n^y, \quad (1)$$

where M^y is the number of one values in the pseudo domain label \mathbf{d}^y . With this aggregated style vector $\mathbf{a}^y \in \mathbb{R}^{1 \times s}$ and content vector \mathbf{c}^x , our networks finally generates the output such that $\hat{\mathbf{y}} = \mathcal{G}(\mathbf{c}^x, \mathbf{a}^y)$, containing a style normalization layer. The detailed network architecture is described in the suppl. material.

3.3. Language-Driven Domain Labeling

In this section, we explain how to obtain the pseudo multi-hot domain label \mathbf{d}^y for an image \mathbf{y} using pretrained vision-language model. Unlike TUNIT [3] that clusters the images from the dataset using the visual features only, we use n -th domain’s textual description as prompt $\mathbf{t}_n \in \mathbb{R}^{1 \times L}$ which is defined as

$$\mathbf{t}_n = [p_1, p_2, \dots, p_L, p_n^{\text{domain}}], \quad (2)$$

which consists of one of candidate domains $\{p_n^{\text{domain}}\}_{n=1}^N$ and a template p_l ($l \in \{1, \dots, L\}$) that is shared with all the candidate domains, where p_l is a vector for a word with the same dimension as p_n^{domain} . For instance, all p_l and p_n^{domain} can be initialized with a template, e.g., “a face with”, and given textual domain descriptions, e.g., “black hair”, respectively. It is utilized to softly cluster the images by measuring the similarity between visual features and language features. In our framework, all the prompts $\mathbf{T} = [\mathbf{t}_n]_{n=1}^N \in \mathbb{R}^{N \times (L+1)}$ are initialized by the human-annotated domain descriptions or keywords, namely *dataset-level* supervision.

Based on this definition, we first extract vision and language features $\mathbf{v}^y = \mathcal{E}_V(\mathbf{y}) \in \mathbb{R}^{1 \times k}$ with k channels, and $[\mathbf{u}_n^t]_{n=1}^N = \mathcal{E}_L(\mathbf{T}) \in \mathbb{R}^{N \times k}$ with k channels from style \mathbf{y} and all the prompts \mathbf{T} , respectively, by using pretrained vision-language model, i.e., CLIP [49]. Note that any vision-language models with vision and language encoders, i.e., ALIGN [24], can be used in our framework. We then measure a similarity $\mathbf{f}^y = [f_n^{\mathbf{y}, \mathbf{t}}]_{n=1}^N \in \mathbb{R}^{N \times 1}$, computed as:

$$f_n^{\mathbf{y}, \mathbf{t}} = \bar{\mathbf{v}}^y \cdot \bar{\mathbf{u}}_n^t, \quad (3)$$

where $\bar{\mathbf{v}} = \mathbf{v} / \|\mathbf{v}\|$ and $\bar{\mathbf{u}} = \mathbf{u} / \|\mathbf{u}\|$.

By utilizing the measured similarity \mathbf{f}^y , we can obtain the multi-hot pseudo domain label \mathbf{d}^y with a simple candidate selecting method, e.g., top- K operation or thresholding. However, with top- K operation, we can just select K number of attributes ignoring the other attributes that would exist in the style image. Moreover, a simple thresholding with a predefined hyperparameter would limit the performance. To overcome these, we present an adaptive thresholding technique based on the observation that the more specified text descriptions, e.g., “a face with blonde hair” make higher similarity with corresponding images than the base prompt as a neutral text description, e.g., “a face with”, that represent each dataset. We set the base prompt \mathbf{p} as the template, meaning that they are updated simultaneously, and thus $\mathbf{u}^{\mathbf{p}} = \mathcal{E}_L(\mathbf{p}) \in \mathbb{R}^{1 \times k}$. We then define the domain label $\mathbf{d}^y = [d_n^y]_{n=1}^N \in \mathbb{R}^{N \times 1}$ as follows:

$$d_n^y = \begin{cases} 1, & \text{if } f_n^{\mathbf{y}, \mathbf{t}} > \bar{\mathbf{v}}^y \cdot \bar{\mathbf{u}}^{\mathbf{p}}, \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

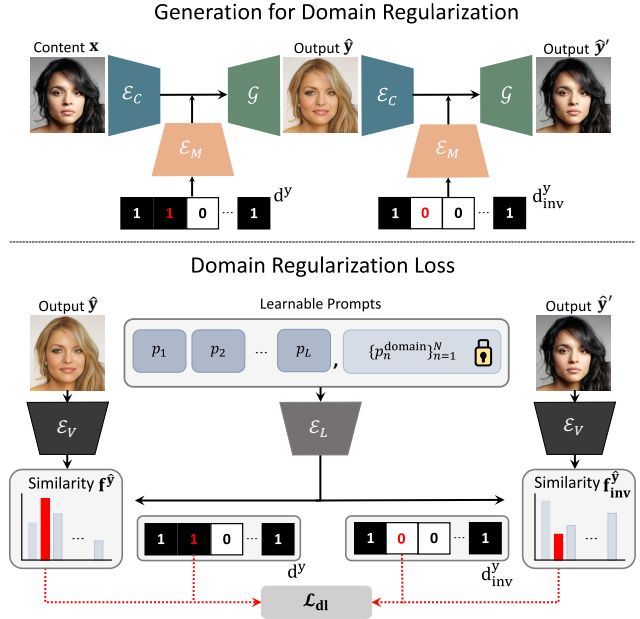


Figure 4. **Illustration of domain regularization loss.** We define a domain regularization loss \mathcal{L}_{d1} utilizing two outputs $\hat{\mathbf{y}}$ and $\hat{\mathbf{y}}'$ that have the opposite styles at n -th domain, and minimize the loss function not to only learn the optimal prompt but also better learn the translation process.

3.4. Prompt Learning

So far, we have discussed the method for language-driven domain labeling and image translation with the pseudo domain label. However, since the pre-defined templates may not be optimal to describe all the images in the dataset, the performance would be noisy and limited. To overcome these, we present a prompt learning technique as shown in Fig. 4.

Similar to [20, 64, 65], we set the prompt \mathbf{t}_n as learnable except p_n^{domain} . In other words, the templates $[p_l]_l$ are set as the learnable parameters in \mathbf{t}_n . By allowing the template to be updated, more optimal prompts can be optimized, which in turn boosts the translation performance. Note that the input domain descriptions should be given in our framework, but they can also be obtained by using a predefined dictionary [1, 49] in a manner that the highly relevant texts to the dataset can be selected based on the similarity between an image and candidate texts.

3.5. Loss Functions

Adversarial Loss. We adopt an adversarial loss to guide the translated images $\hat{\mathbf{y}}$ to arrive at the training distribution of the target domain. Following StarGAN2 [10] and TUNIT [3], we also adopt the multi-domain discriminators, but we use the outputs of multi-domain discriminators weighted

by a multi-hot domain label \mathbf{d}^y :

$$\begin{aligned} \mathcal{L}_{\text{adv}} &= \mathbb{E}_{\mathbf{x}, \mathbf{y}} \sum_{n=1}^N [\log \mathcal{D}_n(\mathbf{y}) d_n^y + \log(1 - \mathcal{D}_n(\mathcal{G}(\mathbf{x}, \mathbf{a}^y)) d_n^y)], \end{aligned} \quad (5)$$

where $\mathcal{D}_n(\cdot)$ denotes n -th discriminator output and d_n^y denotes n -th element of \mathbf{d}^y . It should be noted that if we set \mathbf{d}^y as a one-hot domain label by ground-truth or pseudo-label, this loss function becomes the same adversarial loss in [3, 10].

Domain Regularization Loss. We further present a domain regularization loss that encourages the networks to generate an image $\hat{\mathbf{y}}$ following the domain label of input style image \mathbf{y} more precisely. To this end, perhaps one of the simplest functions is to use the dot product between \mathbf{d}^y and $\mathbf{f}^{\hat{\mathbf{y}}}$ such that $\mathcal{L} = \mathbf{d}^y \cdot \mathbf{f}^{\hat{\mathbf{y}}}$. Minimizing this, however, can induce erroneous solutions, since considering an inactive domain label, i.e., $d_n^y = 0$, may degrade the performance further. To overcome this, we present a novel loss function. We first make a domain label pair, \mathbf{d}^y and $\mathbf{d}_{\text{inv}}^y(n)$, which has same labels with \mathbf{d}^y except the n -th label by replacing n -th domain label d_n^y to $d_{\text{inv},n}^y(n) = 1 - d_n^y$. We then generate $\hat{\mathbf{y}}' = \mathcal{G}(\mathbf{c}^{\hat{\mathbf{y}}}, \mathbf{a}_{\text{inv}}^{\hat{\mathbf{y}}})$ where $\mathbf{a}_{\text{inv}}^{\hat{\mathbf{y}}}$ is obtained using $\hat{\mathbf{y}}$ and $\mathbf{d}_{\text{inv}}^y(n)$. We then encourage not only $\mathbf{f}^{\hat{\mathbf{y}}}$ to follow \mathbf{d}^y at n -th label, but also $\mathbf{f}^{\hat{\mathbf{y}}_{\text{inv}}}$ to follow $\mathbf{d}_{\text{inv}}^y(n)$ at n -th label, such that

$$\mathcal{L}_{\text{dl}} = \mathcal{H}(d_n^y, f_n^{\hat{\mathbf{y}}}) + \mathcal{H}(d_{\text{inv},n}^y(n), f_n^{\hat{\mathbf{y}}_{\text{inv}}}), \quad (6)$$

where $\mathcal{H}(\cdot, \cdot)$ denotes a binary cross-entropy.

Cycle-Consistency Loss. To make the translated image similar to its original image, which also regularizes the ill-posed translation problem, we adopt a cycle-consistency loss such that

$$\mathcal{L}_{\text{cyc}} = \mathbb{E}_{\mathbf{x}, \mathbf{y}} [\|\mathbf{x} - \mathcal{G}((\hat{\mathbf{y}}), \mathbf{a}^x)\|_1], \quad (7)$$

where $\mathbf{c}^{\hat{\mathbf{y}}}$ denotes the content from $\hat{\mathbf{y}}$, and \mathbf{a}^x denotes the style vector of input \mathbf{x} . By encouraging the generator \mathcal{G} to reconstruct the input image \mathbf{x} with estimated style vector \mathbf{a}^x , \mathcal{G} learns to preserve the original characteristics of \mathbf{x} .

Style Reconstruction Loss. While the generator is able to synthesize realistic images with the losses above, the synthesized results are not guaranteed to be style-consistent with \mathbf{y} . In order to better learn style representation, we compute l_1 loss between the style vector from the translated image and style image such that

$$\mathcal{L}_{\text{sty}} = \mathbb{E}_{\mathbf{x}, \mathbf{y}} [\|\mathbf{s}^y - \mathcal{E}_S(\hat{\mathbf{y}})\|_1]. \quad (8)$$

Style Diversification Loss. To learn the generator to produce diverse images, we employ a diversity sensitive

loss [43, 58] defined as:

$$\mathcal{L}_{\text{ds}} = \mathbb{E}_{\mathbf{x}, \mathbf{y}} [\|\mathcal{G}(\mathbf{x}, \mathcal{E}_M(\mathbf{z}_1)) - \mathcal{G}(\mathbf{x}, \mathcal{E}_M(\mathbf{z}_2))\|_1], \quad (9)$$

where \mathbf{z}_1 and \mathbf{z}_2 are random latent vectors from Gaussian distribution. Note that we only adopt the diversification loss using random vectors, not the style vector \mathbf{a}^y from an image. Also, we utilize the domain label \mathbf{d}^y from the style image \mathbf{y} as a pseudo-label for the latent vector.

Overall Objective. Full loss functions are as follows:

$$\begin{aligned} \mathcal{L}_{\text{total}} &= \lambda_{\text{adv}} \mathcal{L}_{\text{adv}} + \lambda_{\text{dl}} \mathcal{L}_{\text{dl}} + \lambda_{\text{cyc}} \mathcal{L}_{\text{cyc}} \\ &\quad + \lambda_{\text{sty}} \mathcal{L}_{\text{sty}} - \lambda_{\text{ds}} \mathcal{L}_{\text{ds}}, \end{aligned} \quad (10)$$

where λ_{adv} , λ_{dl} , λ_{cyc} , λ_{sty} , and λ_{ds} are hyper-parameters.

4. Experiments

Implementation Details We adopted StarGAN2 [10] as the baseline architecture for content, style, mapping encoders, generator, and discriminator. Additional implementation details will be explained in the supplementary material.

Setting up Templates and Candidate Domains. We set the candidate domains by randomly sampling among the domain or attribute names in each dataset. For a fair comparison, we experiment three times with different domains respectively to measure quantitative results. The selected candidate domains are shown in the supplementary materials. We refer to the templates from [5, 12] and slightly modify them to align for each dataset. To improve the accuracy of pseudo labels, we also adopt template augmentation in our baseline model.

Datasets. We conduct the experiments on four standard datasets, including Animal Faces-10, Food-10, CelebA-HQ, and LHQ [7, 35, 41, 54]. In particular, we consider 10 textual domain descriptions among Animal Faces-149 and Food-101, respectively, selected in TUNIT [3], to define the initial prompt of our framework. For CelebA-HQ, we obtain 40 attributes, and sample 10 attributes for three times to define the initial prompt and report the average results. Since LHQ dataset [54] does not provide any attributes, we sampled various target domain descriptions and selected 10 domains for experiments.

Evaluation Metrics. In our experiments, we adopt three quantitative evaluation metrics. First, the mean of class-wise Fréchet Inception Distance (mFID) [17] is used to evaluate how the translated image reflects the input style. In addition, we use Density and Coverage (D&C) [44], which measure the fidelity and diversity of the translated images, respectively. The lower the mFID value, the better quality of the image. As fidelity and diversity get better, the D&C scores are higher or closer to 1. We also measure Acc., namely classification accuracy, between the domain labels with the highest probability and ground-truth domain labels.




Figure 5. Qualitative comparison.

Method	CelebA-HQ [41]		AnimalFaces-10 [35]		Food-10 [7]	
	mFID ↓	D&C ↑	mFID ↓	D&C ↑	mFID ↓	D&C ↑
StarGAN2 [10] (sup.)	32.16	1.22 / 0.44	33.67	1.54 / 0.91	65.03	1.09 / 0.76
Smoothing [40] (sup.)	35.93	1.25 / 0.43	38.93	0.97 / 0.75	61.13	0.96 / 0.68
TUNIT [3] (unsup.)	61.29	0.24 / 0.13	47.70	1.04 / 0.81	52.20	1.08 / 0.88
Kim <i>et al.</i> [29] (unsup.)	41.33	0.60 / 0.24	36.83	1.06 / 0.82	49.34	1.06 / 0.80
LANIT	27.96	0.91 / 0.34	34.11	1.46 / 0.89	48.08	1.24 / 0.86

Table 1. Quantitative comparison on CelebA-HQ [41], Animal Faces-10 [35] and Food-10 [7]. The configurations of StarGAN2 and Smoothing use ground-truth domain labels while TUNIT and Kim *et al.* use pseudo-labels generated from each image. Our LANIT uses only textual domain descriptions.

4.1. Comparisons for Other Models

Quantitative Results. We first report the quantitative comparison of our model with StarGAN2 [10], Smoothing [40] (supervised), TUNIT [3] and Kim *et al.* [29] (unsupervised) on CelebA-HQ, AnimalFaces-10 and Food-10 datasets. Note that StarGAN2 and Smoothing are trained using 10 classes of ground-truth *per-sample* domain labels. Our and truly-unsupervised models use GT labels for calculating quantitative metrics. Tab. 1 shows the quantitative comparison in terms of mFID, D&C. metrics. LANIT shows the best mFID and D&C scores, proving that our model outperforms the state-of-the-art methods in terms of image quality, fidelity and diversity. We also observe that LANIT consistently outperforms unsupervised models by a large margin on all three datasets on Acc. metric. Although our LANIT only utilizes unlabeled datasets with *dataset-level* text supervision, it shows competitive results against

StarGAN2 [10] and Smoothing [40], which require ground-truth per-sample domain labels. Note that we do not provide quantitative results on LHQ since it does not provide ground-truth labels.

Qualitative Results. We show visual results for Tab. 1 in Fig. 5. We observe that the quality differences between our LANIT and TUNIT [3] or Kim *et al.* [29] are significant. The images generated by the other methods tend to contain artifacts and show limited performance in reflecting diverse attributes from style images, while our method yields stable performance in terms of style representation and semantic consistency. Especially, we show that our model is highly capable of capturing multiple subtle attributes, such as “smiling” or “straight hair”, while TUNIT [3] and Kim *et al.* [29] often fail. Interestingly, our LANIT shows competitive results compared to StarGAN2 [10] and Smoothing [40] which are supervised methods, even with better

N	Method	AnimalFaces-10 [35]		CelebA-HQ [41]	
		mFID ↓	D&C ↑	mFID ↓	D&C ↑
4	TUNIT	77.7	0.88 / 0.74	61.5	0.24 / 0.12
	LANIT	71.6	1.35 / 0.46	49.3	0.33 / 0.14
7	TUNIT	62.7	1.02 / 0.73	54.7	0.33 / 0.16
	LANIT	49.9	1.47 / 0.66	43.2	0.44 / 0.19
10	TUNIT	47.7	1.04 / 0.81	61.3	0.24 / 0.13
	LANIT	34.1	1.46 / 0.89	27.9	0.91 / 0.34
13	TUNIT	56.8	0.99 / 0.72	98.9	0.08 / 0.03
	LANIT	30.1	1.43 / 0.85	34.8	0.58 / 0.21
16	TUNIT	54.1	1.09 / 0.78	127.7	0.04 / 0.02
	LANIT	35.8	1.49 / 0.82	27.9	0.76 / 0.23

Table 2. Quantitative comparison of LANIT with TUNIT [3] by varying the number of domains.

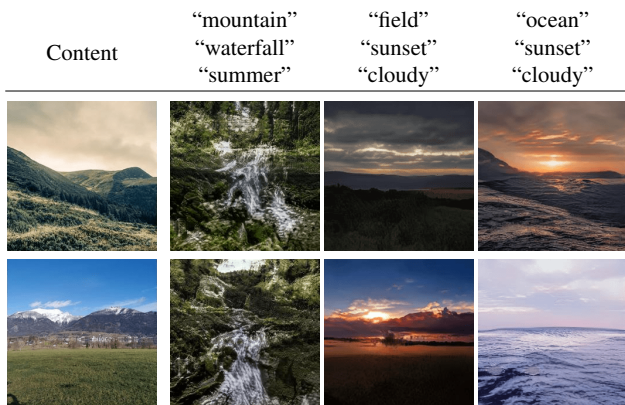


Figure 6. Latent-guided diverse image synthesis results by our LANIT on LHQ [54]. Target domain text descriptions are shown above each image.

style preservation. We also show domain-to-domain diverse image synthesis results in Fig. 6, which is guided by target domain labels. Thanks to our proposed prompt learning, the mapping network and style encoder can faithfully produce the style vectors reflecting the target multiple styles.

Comparisons to other CLIP-based methods. We compare our LANIT with CLIP-based text-guided manipulation models, DiffusionCLIP [28] and StyleCLIP [47], on CelebA-HQ in Fig. 7. In the results, DiffusionCLIP and StyleCLIP show a limited capacity to synthesize realistic results with multiple attributes, while our LANIT generates high-quality images, faithfully representing each domain description, regardless of the number of domain descriptions. The reason is that our model focuses on learning style manifold during training from each target domain defined by textual descriptions, while others are designed to just fine-tune the pre-trained generative models on each instance independently. We observe that LANIT also works on a flexible number of K during inference time, which demonstrates the controllability of our model.

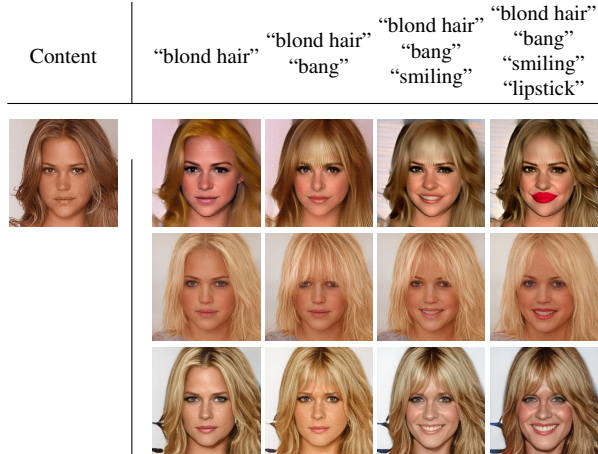


Figure 7. Additional qualitative results by (from top to bottom) DiffusionCLIP [28], StyleCLIP [47], and our LANIT.

Datasets	Default		Dictionary	
	mFID ↓	D&C ↑	mFID ↓	D&C ↑
CelebA-HQ [15]	27.96	0.91 / 0.34	28.34	0.77 / 0.23
AnimalFaces-10 [35]	34.11	1.46 / 0.89	40.48	1.01 / 0.78
Food-10 [7]	48.08	1.24 / 0.86	49.50	1.17 / 0.81

Table 3. Quantitative results with pre-defined dictionary.

4.2. Ablation Study and Analysis

Number of Domain Descriptions N . In Tab. 2, We validate the effects of the number of candidate domains N compared to TUNIT [3]. Our results show that a different number of N brings minor changes to our model than TUNIT [3], which justifies the robustness of our framework. In animal faces-10, LANIT shows the best performance when there exist 10 candidate domains, whereas in CelebA-HQ, there is little difference in the FID score when N is over 10.

Predefined Dictionary. We also evaluate the quantitative results when using a predefined dictionary instead of exploiting specific domain descriptions that we manually provide. For this experiment, we regard the whole class names in each dataset as a dictionary. For sampling candidate domains, we first calculated similarity values on the whole images and target domains, and then only averaged the similarity values assigned to the target domains in the dictionary. Then, we ordered average values in descending order and selected 10 domains. While our LANIT shows a slight drop in both mFID and F1 scores as shown in the Tab. 3 for such a setting in all the datasets, it still records competitive performance compared to the default setting.

Adaptive Thresholding with Base Prompt. We study the effectiveness of the adaptive thresholding technique using base prompt compared to the simple top- K method in Tab. 4. We find that the proposed technique brings a large improvement in performance. This indicates that it enables the model to mitigate the influence of ambiguous or noisy




Figure 8. Qualitative results for ablation study.

Method	CelebA-HQ [41]		
	mFID ↓	D&C ↑	F1 ↑
A LANIT (Top-1)	49.65	0.56 / 0.32	0.347
B LANIT (Top-3)	41.68	0.68 / 0.30	0.564
C Baseline	33.79	0.70 / 0.26	0.613
D + DL Loss	29.21	0.86 / 0.32	0.632
E + Prompt Learning	27.96	0.91 / 0.34	0.670

Table 4. Ablation Study.

Method	CelebA-HQ [41]		
	mFID ↓	D&C ↑	F1 ↑
End-to-End	30.22	0.81 / 0.27	0.583
Stage-wise	27.96	0.91 / 0.34	0.670

Table 5. Ablation Study for PL.

data. Additionally, we evaluate our model with single or multiple activated top- K attributes. We can see that the result of Top-3 significantly surpasses Top-1. This demonstrates that an image can be described more precisely when given with multiple keywords or text descriptions, while most prior works including TUNIT, Kim *et al.* and StarGAN2 assume to assign a single label for each image. In addition, we evaluate the qualitative results in Fig. 8 corresponding to the Tab. 4. It shows that our LANIT captures more diverse attributes rather than Top-1 and Top-3 models.

Domain Regularization Loss. As shown in Tab. 4 (D and E), our LANIT improves performance Domain Regularization(DL) loss compared to the baseline on CelebA-HQ. Especially, D&C metrics are significantly increased by adding our proposed domain regularization loss. In Fig. 8, the generated image has the style of reference image more faithfully by using domain regularization loss.

Prompt Learning. In order to evaluate the effectiveness of prompt learning, we compare the estimated pseudo-domain labels with and without the proposed prompt learning, where the latter can be regarded as an offline clustering. It shows that with prompt learning, the probability of an image/text prompt pair is more accurately estimated.

Since prompt learning helps reduce label noise by refining the prompt, the result without prompt learning shows limited clustering performance, which is verified in Tab. 4 (E). In addition, we further analyze the effect of when to start prompt learning Tab. 5. The stage-wise learning, where we start prompt learning after convergence of the generative model, shows better performance than end-to-end learning with prompt learning. The reason is that the well-generated image corresponding to a given prompt has a critical role in prompt learning. Note that the base prompt is also updated with prompt learning simultaneously.

5. Conclusion

In this paper, we propose LANguage-driven Image-to-image Translation framework for Unlabeled Data (LANIT) that leverages *dataset-level* supervision. By using textual domain descriptions as dataset-level supervision, our model not only mitigates the heavy dependency on per-sample supervision but also considers multi-hot domain labels for an image. The proposed labeling method allows users to explicitly define the semantic meaning of the target domain, which was not possible in the previous truly-unsupervised I2I framework. In addition, as we observe that a naive adoption of a pre-trained vision-language model has limited controllability and even drops the performance of the model, we present some techniques to overcome such limitations: adaptive thresholding, domain regularization loss, and prompt learning. Our experiments show that our LANIT performs better than truly-unsupervised methods in controllability on the application. Furthermore, the quantitative results of our framework have scored competitive or even better values than the existing image-to-image translation models that require per-sample supervision.

Acknowledgements. This research was supported by the MSIT, Korea (IITP-2023-2020-0-01819, ICT Creative Consilience program, No. 2020-0-00368, A Neural-Symbolic Model for Knowledge Acquisition and Inference Techniques), and National Research Foundation of Korea (NRF2021R1C1C1006897).

References

- [1] Rameen Abdal, Peihao Zhu, John Femiani, Niloy J Mitra, and Peter Wonka. Clip2stylegan: Unsupervised extraction of stylegan edit directions. *arXiv preprint arXiv:2112.05219*, 2021.
- [2] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18208–18218, 2022.
- [3] Kyungjune Baek, Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Hyunjun Shim. Rethinking the truly unsupervised image-to-image translation. In *ICCV*, pages 14154–14163, 2021.
- [4] Hyojin Bahng, Sunghyo Chung, Seungjoo Yoo, and Jaegul Choo. Exploring unlabeled faces for novel attribute discovery. In *CVPR*, pages 5821–5830, 2020.
- [5] Omer Bar-Tal, Dolev Ofri-Amar, Rafail Fridman, Yoni Kasten, and Tali Dekel. Text2live: Text-driven layered image and video editing. In *European Conference on Computer Vision*, pages 707–723. Springer, 2022.
- [6] Deblina Bhattacharjee, Seungryong Kim, Guillaume Vizier, and Mathieu Salzmann. Dunit: Detection-based unsupervised image-to-image translation. In *CVPR*, pages 4787–4796, 2020.
- [7] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *ECCV*, pages 446–461. Springer, 2014.
- [8] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 33:1877–1901, 2020.
- [9] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *CVPR*, pages 8789–8797, 2018.
- [10] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *CVPR*, 2020.
- [11] Aviv Gabbay and Yedid Hoshen. Scaling-up disentanglement for image translation. *arXiv preprint arXiv:2103.14017*, 2021.
- [12] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022.
- [13] Chunjiang Ge, Rui Huang, Mixue Xie, Zihang Lai, Shiji Song, Shuang Li, and Gao Huang. Domain adaptation via prompt learning. *arXiv preprint arXiv:2202.06687*, 2022.
- [14] Abel Gonzalez-Garcia, Joost Van De Weijer, and Yoshua Bengio. Image-to-image translation for cross-domain disentanglement. *arXiv preprint arXiv:1805.09730*, 2018.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [16] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022.
- [17] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, pages 6626–6637, 2017.
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *arXiv preprint arxiv:2006.11239*, 2020.
- [19] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, pages 1989–1998. PMLR, 2018.
- [20] Tony Huang, Jack Chu, and Fangyun Wei. Unsupervised prompt learning for vision-language models. *arXiv preprint arXiv:2204.03649*, 2022.
- [21] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, pages 172–189, 2018.
- [22] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, pages 1125–1134, 2017.
- [23] Somi Jeong, Youngjung Kim, Eungbean Lee, and Kwanghoon Sohn. Memory-guided unsupervised image-to-image translation. In *CVPR*, pages 6558–6567, 2021.
- [24] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, pages 4904–4916. PMLR, 2021.
- [25] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *European Conference on Computer Vision*, pages 105–124. Springer, 2022.
- [26] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410, 2019.
- [27] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020.
- [28] Gwanghyun Kim and Jong Chul Ye. Diffusionclip: Text-guided image manipulation using diffusion models. 2021.
- [29] Kunhee Kim, Sanghun Park, Eunyeong Jeon, Taehun Kim, and Daijin Kim. A style-aware discriminator for controllable image translation. In *CVPR*, 2022.
- [30] Soohyun Kim, Jongbeom Baek, Jihye Park, Gyeongnyeong Kim, and Seungryong Kim. Instaformer: Instance-aware image-to-image translation with transformer. In *CVPR*, pages 18321–18331, 2022.
- [31] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *ECCV*, pages 35–51, 2018.
- [32] Wanhua Li, Xiaoke Huang, Zheng Zhu, Yansong Tang, Xiu Li, Jie Zhou, and Jiwen Lu. Ordinalclip: Learning

- rank prompts for language-guided ordinal regression. *arXiv preprint arXiv:2206.02338*, 2022.
- [33] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021.
- [34] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NeurIPS*, pages 700–708, 2017.
- [35] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. Few-shot unsupervised image-to-image translation. In *arxiv*, 2019.
- [36] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*, 2021.
- [37] Steven Liu, Tongzhou Wang, David Bau, Jun-Yan Zhu, and Antonio Torralba. Diverse image generation via self-conditioned gans. In *CVPR*, pages 14286–14295, 2020.
- [38] Xingchao Liu, Chengyue Gong, Lemeng Wu, Shujian Zhang, Hao Su, and Qiang Liu. Fusedream: Training-free text-to-image generation with improved clip+ gan space optimization. *arXiv preprint arXiv:2112.01573*, 2021.
- [39] Xihui Liu, Dong Huk Park, Samaneh Azadi, Gong Zhang, Arman Chopikyan, Yuxiao Hu, Humphrey Shi, Anna Rohrbach, and Trevor Darrell. More control for free! image synthesis with semantic diffusion guidance. *arXiv preprint arXiv:2112.05744*, 2021.
- [40] Yahui Liu, Enver Sangineto, Yajing Chen, Linchao Bao, Haoxian Zhang, Nicu Sebe, Bruno Lepri, Wei Wang, and Marco De Nadai. Smoothing the disentangled latent style space for unsupervised image-to-image translation. In *CVPR*, pages 10785–10794, 2021.
- [41] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, pages 3730–3738, 2015.
- [42] Mario Lučić, Michael Tschannen, Marvin Ritter, Xiaohua Zhai, Olivier Bachem, and Sylvain Gelly. High-fidelity image generation with fewer labels. In *ICML*, pages 4183–4192. PMLR, 2019.
- [43] Qi Mao, Hsin-Ying Lee, Hung-Yu Tseng, Siwei Ma, and Ming-Hsuan Yang. Mode seeking generative adversarial networks for diverse image synthesis. In *CVPR*, pages 1429–1437, 2019.
- [44] Muhammad Ferjad Naeem, Seong Joon Oh, Youngjung Uh, Yunjey Choi, and Jaejun Yoo. Reliable fidelity and diversity metrics for generative models. In *ICML*, pages 7176–7185. PMLR, 2020.
- [45] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.
- [46] Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *ECCV*, pages 319–345. Springer, 2020.
- [47] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *ICCV*, pages 2085–2094, 2021.
- [48] Fabio Petroni, Tim Rocktäschel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, Alexander H Miller, and Sebastian Riedel. Language models as knowledge bases? *arXiv preprint arXiv:1909.01066*, 2019.
- [49] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021.
- [50] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *arXiv preprint arXiv:2208.12242*, 2022.
- [51] Kuniaki Saito, Kate Saenko, and Ming-Yu Liu. Cocommit: Few-shot unsupervised image translation with a content conditioned style encoder. In *ECCV*, pages 382–398. Springer, 2020.
- [52] Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980*, 2020.
- [53] Manli Shu, Weili Nie, De-An Huang, Zhiding Yu, Tom Goldstein, Anima Anandkumar, and Chaowei Xiao. Test-time prompt tuning for zero-shot generalization in vision-language models. *arXiv preprint arXiv:2209.07511*, 2022.
- [54] Ivan Skorokhodov, Grigori Sotnikov, and Mohamed Elhoseiny. Aligning latent and image spaces to connect the unconnectable. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14144–14153, 2021.
- [55] Yaxing Wang, Salman Khan, Abel Gonzalez-Garcia, Joost van de Weijer, and Fahad Shahbaz Khan. Semi-supervised learning for few-shot image-to-image translation. In *CVPR*, pages 4453–4462, 2020.
- [56] Tianyi Wei, Dongdong Chen, Wenbo Zhou, Jing Liao, Zhen-tao Tan, Lu Yuan, Weiming Zhang, and Nenghai Yu. Hairclip: Design your hair by text and reference image. *arXiv preprint arXiv:2112.05142*, 2021.
- [57] Po-Wei Wu, Yu-Jing Lin, Che-Han Chang, Edward Y Chang, and Shih-Wei Liao. Relgan: Multi-domain image-to-image translation via relative attributes. In *ICCV*, pages 5914–5922, 2019.
- [58] Dingdong Yang, Seunghoon Hong, Yunseok Jang, Tianchen Zhao, and Honglak Lee. Diversity-sensitive conditional generative adversarial networks. In *ICLR*, 2019.
- [59] Yuan Yao, Ao Zhang, Zhengyan Zhang, Zhiyuan Liu, Tat-Seng Chua, and Maosong Sun. Cpt: Colorful prompt tuning for pre-trained vision-language models. *arXiv preprint arXiv:2109.11797*, 2021.
- [60] Zili Yi, Hao Zhang, Ping Tan, and Minglun Gong. Dualgan: Unsupervised dual learning for image-to-image translation. In *ICCV*, pages 2849–2857, 2017.
- [61] Pan Zhang, Bo Zhang, Dong Chen, Lu Yuan, and Fang Wen. Cross-domain correspondence learning for exemplar-based image translation. In *CVPR*, pages 5143–5153, 2020.

- [62] Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. The spatially-correlative loss for various image translation tasks. In *CVPR*, pages 16407–16417, 2021.
- [63] Chunting Zhou, Junxian He, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Prompt consistency for zero-shot task generalization. *arXiv preprint arXiv:2205.00049*, 2022.
- [64] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *arXiv preprint arXiv:2109.01134*, 2021.
- [65] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. *arXiv preprint arXiv:2203.05557*, 2022.
- [66] Xingran Zhou, Bo Zhang, Ting Zhang, Pan Zhang, Jianmin Bao, Dong Chen, Zhongfei Zhang, and Fang Wen. Cocosnet v2: Full-resolution correspondence learning for image translation. In *CVPR*, pages 11465–11475, 2021.
- [67] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.