

Domain-Scalable Unpaired Image Translation via Latent Space Anchoring

Siyu Huang*, Jie An*, Donglai Wei, Zudi Lin, Jiebo Luo, *Fellow, IEEE*, and Hanspeter Pfister, *Fellow, IEEE*

Abstract—Unpaired image-to-image translation (UNIT) aims to map images between two visual domains without paired training data. However, given a UNIT model trained on certain domains, it is difficult for current methods to incorporate new domains because they often need to train the full model on both existing and new domains. To address this problem, we propose a new domain-scalable UNIT method, termed as *latent space anchoring*, which can be efficiently extended to new visual domains and does not need to fine-tune encoders and decoders of existing domains. Our method anchors images of different domains to the same latent space of frozen GANs by learning lightweight encoder and regressor models to reconstruct single-domain images. In the inference phase, the learned encoders and decoders of different domains can be arbitrarily combined to translate images between any two domains without fine-tuning. Experiments on various datasets show that the proposed method achieves superior performance on both standard and domain-scalable UNIT tasks in comparison with the state-of-the-art methods.

Index Terms—Unsupervised image-to-image translation, multi-domain image translation, semantic structure alignment, GANs prior.

1 INTRODUCTION

IMAGE-TO-IMAGE translation aims to map images from one visual domain to another (Fig. 1a), where visual domains can be different image modalities (e.g., RGB \leftrightarrow sketch), styles (e.g., winter \leftrightarrow summer), or objects (e.g., horse \leftrightarrow zebra). Due to the scarcity of paired data, unsupervised image-to-image translation (UNIT) methods [1], [2], [3] have been proposed to learn a bi-directional mapping between image domains (Fig. 1b). However, it is challenging to extend UNIT into more than two domains. Multi-domain UNIT methods are thus proposed to avoid the complexity of learning translators for all pairs of domains, which maintain a shared latent space for all available domains (Fig. 1c) while an encoder and a decoder of each domain are trained to map images to the shared feature space and invert features back to images, respectively.

One major weakness of existing multi-domain UNIT models is the lack of scalability, *i.e.*, once the model is trained on certain domains, it is difficult to extend the model to unseen domains. The reason behind this is that the shared latent space of existing methods corresponds to the domains where the training data come from. When adding new domains, the shared feature space would be changed. Therefore, one often needs to fine-tune all existing encoders and decoders to make it remap between images and the new shared latent space (Fig. 2a).

In this work, we present a domain-scalable UNIT method, named *latent space anchoring*, which enables an im-

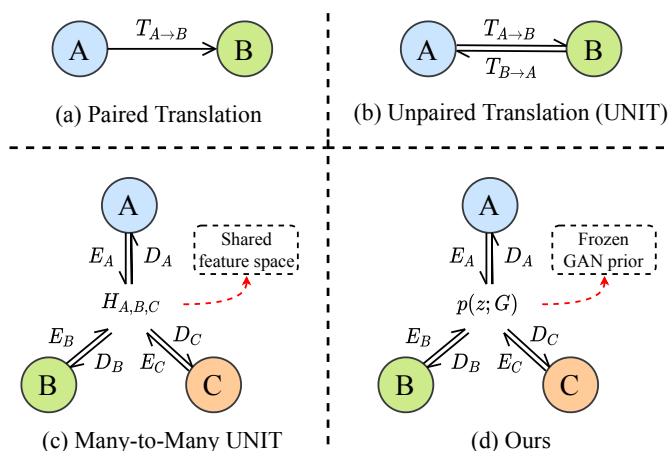


Fig. 1: An illustration of the image-to-image translation paradigms. (a) Dual-domain paired image translation, e.g., Pix2Pix [4] and SPADE [5], learns a one-way mapping $T_{A \rightarrow B}$ from domain A to domain B . (b) Dual-domain unpaired image translation, e.g., CycleGAN [1] and UNIT [2], learns a bi-directional mapping $T_{A \rightarrow B}$ and $T_{B \rightarrow A}$ between two domains. (c) Many-to-many unpaired image translation, e.g., StarGAN [1] and DRIT++ [6], learns mappings between any two domains via a shared feature space H , encoders $E_{\{A,B,C\}}$, and decoders $D_{\{A,B,C\}}$. (d) Our latent space anchoring method learns mappings between any two domains via the latent space $p(z)$ of a frozen GAN G .

- *: *Equal contribution*
- Siyu Huang, Zudi Lin, and Hanspeter Pfister are with the John A. Paulson School of Engineering and Applied Sciences, Harvard University, Boston, MA 02134. E-mail: {huang.lin,zudi.pfister}@seas.harvard.edu
- Jie An and Jiebo Luo are with the Department of Computer Science, University of Rochester, Rochester, NY 14627. E-mail: {jan6,jluo}@cs.rochester.edu
- Donglai Wei is with the Computer Science Department, Boston College, Boston, MA 02116. E-mail: donglai.wei@bc.edu

Manuscript received August 4, 2022.

proved flexibility to extend into new domains without fine-tuning on existing domains. It encodes images from each domain to the latent space of a pretrained GAN via an encoder and aligns the perceptual structure of corresponding images by the reconstruction task via a decoder. As shown in Fig. 2b, to add a new domain X into our model, one only needs to train an encoder and a decoder on the data of a

TABLE 1: A summarization of image-to-image translation tasks and methods. $\{\cdot\}$ denotes the data used in a single training. (\cdot, \cdot) denotes the paired data between domains.

Image Translation Task	Representative Methods	Training Data	Training Data	Inference
		Fixed Domains	Adding a Domain X	
Paired dual-domain I2I	Pix2Pix [4], SPADE [5]	$\{(A, B)\}$	-	$A \rightarrow B$
Unpaired dual-domain I2I	CycleGAN [1], UNIT [2]	$\{A, B\}$	-	$A \leftrightarrow B$
Unpaired multi-domain I2I	StarGAN [7], [8], DRIT++ [6]	$\{A, B, C\}$	$\{A, B, C, X\}$	$A, B, C, X \leftrightarrow A, B, C, X$
	Latent Anchor (Ours)	$\{A\}, \{B\}, \{C\}$	$\{X\}$	$A, B, C, X \leftrightarrow A, B, C, X$

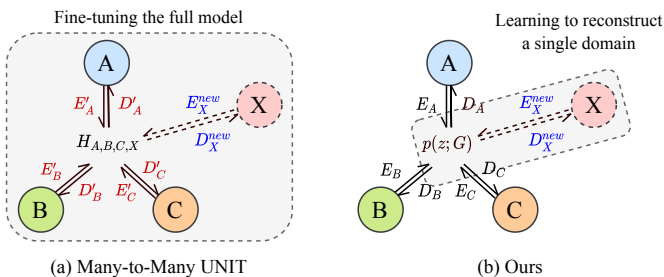


Fig. 2: Adding a new domain X to a learned UNIT model. **Red**: Models that need to be fine-tuned. **Blue**: Models that need to be trained from scratch. **(a)** Existing methods usually need to fine-tune the full model on data of all available image domains. **(b)** Our method only needs to learn an encoder E_X^{new} and a decoder D_X^{new} on data of the new domain X , then images of the domain X can be translated from/to previous domains via the latent space of the generator.

new domain, which are lightweight networks that embed X 's images to the latent space of GANs and reconstruct the images from the latent space. Because the GAN latent space does not depend on specific visual domains, we do not need to fine-tune all existing encoders and decoders when extending the model into new domains. Once the domain X has been incorporated into the existing UNIT model, images of the domain X can be translated to any other existing domains and vice versa, since one can use the encoder of the source domain and the regressor of the target domain to translate images between any two arbitrary domains. The intuition and realization of our latent space anchoring method will be detailed in Sec. 3.

In the experiments, we benchmark our method on various image domains (*e.g.*, human face images, animal face images, natural images, semantic segmentation masks, sketch drawings, and landmark maps), two image-to-image translation settings (*e.g.*, the domain-fixed UNIT setting and the domain-scalable UNIT setting), and two GANs generator backbones (*e.g.*, StyleGAN2 [9] and BigGAN-deep [10]). In quantitative and qualitative comparison with the state-of-the-art methods, our method achieves impressive synthesis performances on challenging multi-domain UNIT tasks characterized by large domain gaps. This is a significant improvement over existing UNIT methods that usually only handle small domain gaps, *i.e.*, slight changes of colors, weather conditions, or facial attributes [1], [7]. User study results also validate the superior visual quality and good structure consistency of the images produced by our method. Ablation studies on the learning objectives and

network architectures further reveal the effects of different pipeline design choices. In summary, our work makes the following contributions:

- We propose a domain-scalable challenge of UNIT methods — to give the models an improved flexibility to extend into new domains without training the full model on the data of all available domains.
- We present a novel UNIT method named latent space anchoring to tackle the challenging task. It aligns the semantic structures of multi-domain images by reconstructing single-domain images from the latent space of frozen GANs.
- We conduct comprehensive experiments on various image domains, UNIT tasks, and different powerful GANs backbones. Qualitative and quantitative comparisons with the state-of-the-art methods, as well as the user study, validate the effectiveness of the proposed method.

2 RELATED WORKS

2.1 Unsupervised Image Translation (UNIT)

The unpaired/unsupervised image translation task aims to translate unpaired images between two domains [3], [11], [12], [13], [14], [15], [16], [17], [18]. One popular approach is based on the cyclic consistency, *e.g.*, CycleGAN [1], DiscoGAN [19], and DualGAN [20], which enforces synthesized images to be faithfully translated back to their original domain. CUT [3] further learns patch-wise correspondence between input and synthesized image patches via contrastive learning based on the CycleGAN [1] scheme. Another approach, *e.g.*, UNIT [2], embeds images of different domains to one shared latent space. In this work, we use the latent space of a pretrained GAN as the intermediate anchor and independently align domains to the anchor to achieve the translation between domains.

2.2 UNIT of Multiple Domains

To efficiently translate images among more than two domains, multi-domain image translation methods [21], [22], [23], *e.g.*, StarGAN [7], build unified models to learn one-to-one mappings between the shared latent space and input domains. To sample multimodal translation results between domains, MUNIT [24] and DRIT [25] decompose the latent space into a shared content space and an unshared style space. Recent works [6], [26], [27], [28], [29], [30], [31], [32], [33] combines the multi-domain and multimodal UNIT. For more details, we refer readers to [34].

Despite the progress, it remains hard to efficiently update an existing model with new domains, which often

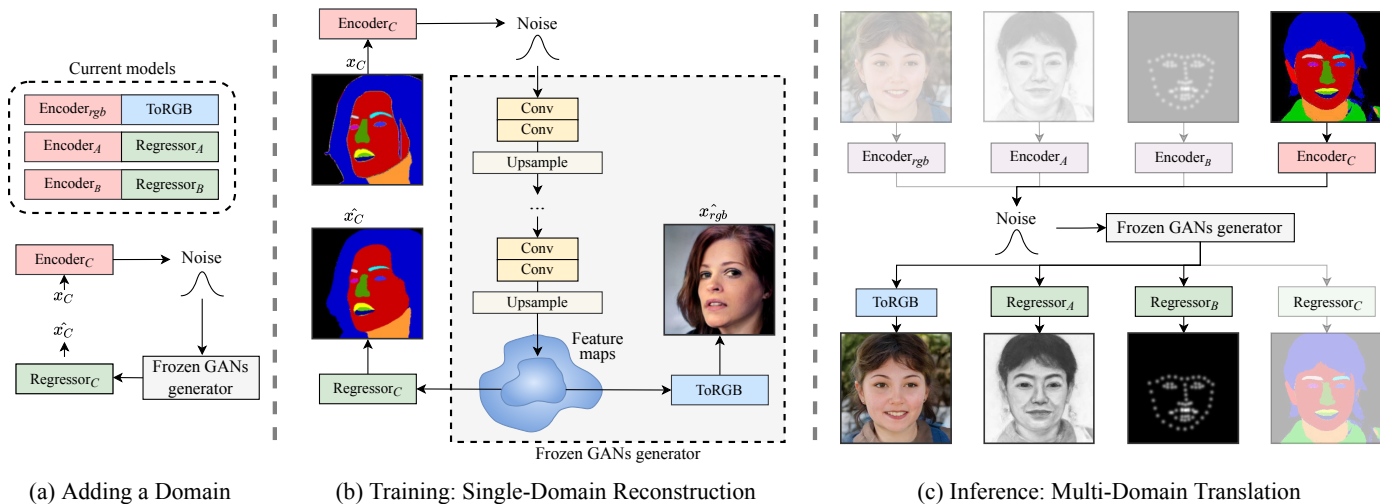


Fig. 3: Our framework for domain-scalable UNIT. **(a)** Adding a new domain C to the current models is easy that one only needs to train an $Encoder_C$ and a $Regressor_C$ on the new domain. **(b)** The training is performed on the single-domain data. We use a lightweight $Encoder_C$ to map the visual domain to GAN’s latent space (Sec. 3.2), then use a lightweight $Regressor_C$ to reconstruct input images from GAN’s latent space. We show that it is equivalent to aligning the perceptual structures of images of different domains. **(c)** During inference, we can use the trained encoder and regressor to perform translation between any two domains, as well as generate well-aligned multi-domain images by sampling from the GAN’s latent space.

involves training or fine-tuning all previous models. To avoid such fine-tuning, ComboGAN [35] trains a separate encoder and decoder for the new domain and aligns its latent space to existing domains via cyclic consistency loss. However, it still requires images from existing domains during training. In contrast, our proposed method only needs the new-domain images. We briefly summarize different image translation tasks and corresponding representative methods in Table 1. Distinct to the existing image translation methods, our method is trained on images of single domains based on the latent space anchoring algorithm. In addition, it is domain-scalable, *i.e.*, it can be efficiently scalable to unseen domains with much fewer training efforts.

2.3 GANs as Image Priors

General Applications. Many existing works have employed pretrained GANs as image priors to various downstream tasks. One approach, *e.g.*, GAN-Dissection [36] enables the manipulation of the semantics of images generated by GANs by discovering their relationship with feature representations. Another approach uses GAN inversion methods [37], [38], [39] to embed images back into the latent space of pretrained GANs, leading to applications in image editing [40], [41], [42], image restoration [38], [43], style transfer [38], [44], and interpreting the latent directions [45], [46]. The proposed method is motivated by the observation that feature maps of pre-trained GAN generators contain rich and semantic structure information. A similar idea has been studied by [47] and [48], where [47] uses pre-trained GAN features for few-shot image annotation generation and [48] utilizes GAN prior to perform unsupervised semantic segmentation. [47], [48] can only perform pixel-aligned translation like image-to-segmentation and cannot work well on domains with structural changes. In this work,

we significantly expand the application scope of this idea to domain-scalable image translation by anchoring images from different visual domains to a shared GANs prior.

GAN Priors to UNIT. More relevant to our work, a few works have exploited pretrained GANs for unidirectional UNIT tasks [49], [50], *i.e.*, GAN-to-Any UNIT. FreezeG [51] freezes the early layers of GANs and optimizes the other layers to generate images of target domains in adversarial training. It only supports mappings from the domain of pretrained GANs to other visual domains. In addition, it does not guarantee a good perceptual structure alignment between translated images (see Fig. 13). For example, if the target domain only has frontal faces, the fine-tuned GAN can hardly generate side-view faces regardless of input images. Instead, we propose a latent space anchoring algorithm that effectively aligns the perceptual structure of images from different domains, enabling a domain-scalable UNIT despite the limitation of the domain dataset.

3 METHOD

3.1 Overview

In this work, we study a new challenging UNIT task, which aims at efficiently incorporating new visual domains to a learned UNIT model. The key idea of the proposed *latent space anchoring* method is to use the latent space of a pretrained GAN as the *anchor*, where each visual domain is aligned independently during training (Fig. 3a). Then during inference, we can achieve a domain-scalable UNIT by encoding any available domain to the GAN’s latent space, then decode the latent vector to any other domain without fine-tuning (Fig. 3c). To exploit the state-of-the-art GAN models [10], [52], we employ an off-the-shelf generator model G that maps Gaussian noise z to RGB images. In




Fig. 4: The visualization results of StyleGAN2 [9] feature maps. We show a channel-wise visualization of randomly selected feature maps from the layer right before the ToRGB module of StyleGAN2.

general, generator G can be split into the feature generator G_{feat} and the last ToRGB layer which has only one or a few neural network layers, *i.e.*, $G = \text{ToRGB} \circ G_{\text{feat}}$.

To align a new domain A to the latent space of generator G , we train an encoder model E_A and a regressor model R_A jointly. We first apply E_A to encode input image x_A to G 's latent code z_A to generate a visually pleasing RGB image x_{rgb} (Sec. 3.2):

$$z_A = E_A(x_A), \quad f_A = G_{\text{feat}}(z_A), \quad x_{\text{rgb}} = \text{ToRGB}(f_A). \quad (1)$$

Then, we align the perceptual structure of the input image x_A and the generated x_{rgb} by equivalently regressing GAN's feature maps to \hat{x}_A with the regressor model R_A , where

$$\hat{x}_A = R_A(f_A). \quad (2)$$

3.2 Image Encoding

Since there are not any matched pairs x_A and x_{rgb} in UNIT setting, learning to reconstruct x_{rgb} from x_A [38], [44] is infeasible. Existing literature [38], [39], [53] shows that the distance between latent code w and mean latent code \bar{w} can be used to evaluate the quality of images generated by StyleGAN [53]. We found that the regularization on latent code z can make z be close to the normal distribution $\mathcal{N}(0, \mathbf{I})$.

For standard GANs, *e.g.* BigGAN [10], we apply regularization to the encoded latent as

$$\mathcal{L}_{\text{latent}} = \|E_A(x_A)\|_2. \quad (3)$$

For style-based GANs [9], [52], [53], we encode the image into the W^+ latent space, and the objective function becomes $\mathcal{L}_{\text{latent}} = \|E_A(x_A) - \bar{w}\|_2$.

Besides $\mathcal{L}_{\text{latent}}$, the adversarial training between \hat{x}_{rgb} and real image set \mathcal{X}_{rgb} is optionally used to enforce z to be on the latent space of pretrained GANs,

$$\mathcal{L}_{\text{adv}} = \mathbb{E}_{x_{\text{rgb}} \sim \mathcal{X}_{\text{rgb}}} [\log D_A(x_{\text{rgb}})] + \mathbb{E}_{x_A \sim \mathcal{X}_A} [\log(1 - D_A(G(E_A(x_A))))], \quad (4)$$

where D_A is the discriminator for domain A .

3.3 Latent Space Anchoring

Training. To align images from multiple domains in GAN's feature representation space, we propose a *latent space anchoring* method to align each visual domain with the latent space of a pretrained GAN, independently. The training objective for a single domain is

$$\mathcal{L}_{\text{total}} = \lambda_{\text{rec}} \mathcal{L}_{\text{rec}} + \lambda_{\text{latent}} \mathcal{L}_{\text{latent}} + \lambda_{\text{adv}} \mathcal{L}_{\text{adv}}. \quad (5)$$

λ_{rec} , λ_{latent} , and λ_{adv} are weights of loss terms, respectively. $\mathcal{L}_{\text{latent}}$ and \mathcal{L}_{adv} push latent features towards the latent

space of the generator. \mathcal{L}_{rec} is the reconstruction loss¹ which ensures the regressor R_A reconstructs realistic images of domain A ,

$$\mathcal{L}_{\text{rec}} = \|\hat{x}_A - x_A\|_2^2. \quad (6)$$

With Eq. 5, the perceptual alignment of domain A and pretrained latent space is achieved by the inductive bias of models, *i.e.*, the easiest way for models to reconstruct realistic images of domain A from latent features is to project samples with the same perceptual structure into the same latent feature. More empirical analysis of the three loss terms can be found in Fig. 12. During training, the encoder E_A learns to anchor multi-domain images to the same latent feature space and the regressor R_A learns to recover multi-domain images from the latent feature space, while the pretrained generator G is fixed.

It is of the pivotal importance to find a semantically rich latent space where images from different domains can be anchored to. Fig. 4 shows randomly sampled feature maps right before the ToRGB layer in StyleGAN2 [9], where features of individual channels show rich, aligned, and diverse semantic structure information, *i.e.*, the perceptual structure. Motivated by this, we use feature maps before the ToRGB layer in GAN generators as the shared anchor ground for every domain.

Inference. As illustrated in Fig. 3c, the GANs latent space serves as a hub of all the domains, such that the encoders and regressors of different modalities can be arbitrarily combined to translate images from either input images (*i.e.*, cross-domain image translation) or random noises (*i.e.*, multi-domain image dataset generation). Assume we have trained encoders and regressors for domain A and domain B . When a new domain C comes, we only need to train an encoder E_C and a regressor R_C for domain C , where the training process is identical to that of domain A and domain B . The images of domain C are aligned with those of domains A and B via the generative latent space \mathcal{H} , thus image translation can be performed between any two of these domains.

4 EXPERIMENTS

4.1 Experimental Setup

4.1.1 Datasets

In the experiments, we evaluate the proposed method on several datasets of multi-domain images. The first dataset consists of facial images from four domains:

1. For domains such as landmark maps and sketch drawings shown in Fig. 3c, the MSE loss function is adopted as the reconstruction loss. For domains such as semantic segmentation maps, the cross-entropy loss is adopted.

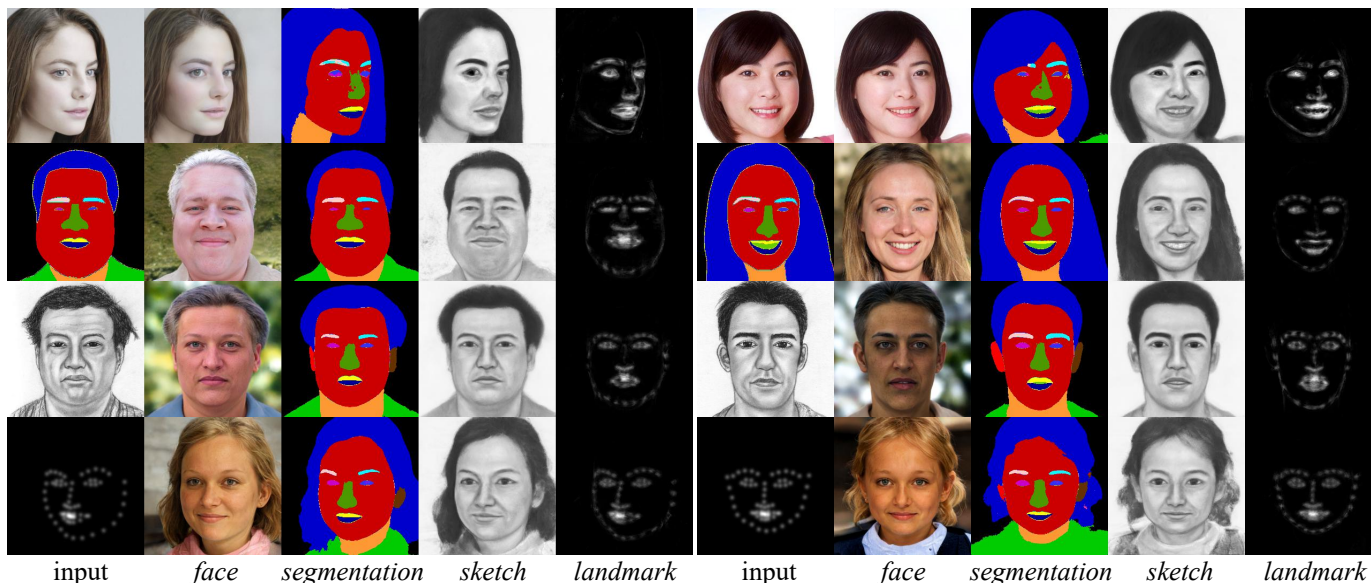


Fig. 5: Unpaired image translation results of our method on four facial domains including RGB image, semantic segmentation map, sketch drawings, and landmark map. Reconstruction results are shown when inputs and outputs are of the same domain.

RGB images, semantic segmentation masks, sketch drawings, and landmark maps. The RGB images come from CelebAMask-HQ [54]. The facial segmentation masks come from CelebAMask-HQ [54], consisting of 30,000 maps of 19 semantic classes. The sketch drawings come from CUHK Face Sketch FERET Database (CUFSF) [55], [56], consisting of 1,194 facial sketch images. The facial landmarks are extracted from FFHQ [53] using the face pose estimator [57] implemented in Dlib toolkit [58]. We aggregate the 68 landmarks to one channel to facilitate the training of image translation baselines. The second dataset consists of human and animal facial images from four domains: human facial images, dog facial images, cat facial images, and wild facial images. The animal images come from the AFHQ dataset [8], and the human faces come from the CelebAMask-HQ dataset [54]. The third dataset consists of natural images from two domains: ImageNet RGB images [59] and foreground object segmentation masks that are extracted by an off-the-shelf method [60], [61].

4.1.2 Network and learning details

Our encoder network architecture follows the pSp encoder [39] which is based on a feature pyramid over the ResNet backbone [62]. The StyleGAN2 [9] and the BigGAN-deep [10] are used as the pretrained generators for face and ImageNet image translation, respectively. The StyleGAN2 [9] is pretrained on CelebAMask-HQ [54], while the BigGAN-deep [10] is pretrained on ImageNet [59]. We embed images into either 128-dimensional z space of BigGAN-deep or 18 W^+ space of StyleGAN2. The regressor is a 6-layer fully convolutional network with feature map channels of (128, 64, 64, 32, 32). ReLU activation and batch normalization are adopted after every convolutional layer except the last layer.

For all experiments, our models are trained with the Adam [63] optimizer using 200k iterations, a learning rate of $1e^{-4}$, and a batch size of 4. For the training objective,

$\lambda_{\text{latent}} = 0.005$, $\lambda_{\text{adv}} = 0.01$, $\lambda_{\text{rec}} = 1$ for the segmentation domain, and $\lambda_{\text{rec}} = 10$ for the sketch and landmark domains. For all translation tasks, we follow the original training/test split. Random 80%/20% split is adopted if the original split is not provided. All images are resized to 256×256 .

4.1.3 Evaluation metrics

We compute FID [64] to evaluate the quality of translated images. All of the paired visual domains, except sketch-to-face, have pairwise annotations. We compute LPIPS [65] between translations and ground truth images to evaluate the correspondences learned by UNIT methods.

4.2 Method Results

4.2.1 UNIT of RGB, segmentation, sketch, and landmark

Fig. 5 shows several examples of facial image translation by our method. The first column shows the input image, and the rest columns show the translations to four domains. We show the reconstruction results if the input and output domains are identical. Fig. 5 demonstrates that our method achieves robust translations between any two domains with diverse facial appearance details, genders, and head poses.

4.2.2 UNIT of cat, dog, wild, and human face

Fig. 6 shows examples of translation between domains of cat, dog, wild, and human face with a StylGAN2 backbone pretrained on the CelebAMask-HQ dataset [54]. Fig. 6 shows that our method achieves faithful AFHQ \leftrightarrow CelebAMask-HQ translations. The shapes, face appearances, and expressions of translations well follow those of the inputs. The color of human skin is bound with the color of fur of the animal face. In addition, the 4th example in Fig. 6 shows that the head pose (e.g., orientation) of translation is consistent with the pose of the input.




Fig. 6: Unpaired image translation results delivered by our method between domains of CelebA-HQ [54], AFHQ-dog [8], AFHQ-cat [8], and AFHQ-wild [8].



Fig. 7: Unpaired facial attribute translation results of our method on the CelebAMask-HQ dataset [54].



Fig. 8: Unpaired image translation between foreground segmentation masks and ImageNet images. The BigGAN-deep [10] model is employed as the generator backbone.

TABLE 2: Quantitative results for the fixed-domain UNIT setting on both facial and ImageNet images.

Method	Facial Images						ImageNet Images					
	seg→face		sketch→face		landmark→face		seg→cat		seg→dog		seg→plane	
	FID↓	LPIPS↓	FID↓	FID↓	LPIPS↓	FID↓	LPIPS↓	FID↓	LPIPS↓	FID↓	LPIPS↓	
CycleGAN [1]	132.6	0.550	120.8	328.1	0.607	278.3	0.709	317.1	0.745	172.1	0.716	
CUT [3]	38.9	0.435	87.6	287.6	0.595	342.9	0.699	350.3	0.752	148.3	0.695	
MUNIT [24]	212.1	0.623	266.4	184.0	0.665	339.9	0.820	372.4	0.902	312.7	0.764	
StarGANv2 [8]	53.7	0.532	85.4	99.5	0.557	-	-	-	-	-	-	
DRIT++ [6]	132.2	0.511	107.5	305.4	0.859	-	-	-	-	-	-	
Latent Anchor (ours)	83.8	0.442	84.3	88.1	0.510	77.9	0.661	79.7	0.654	44.3	0.713	

4.2.3 UNIT of facial attributes

Fig. 7 shows the results of our method for facial attribute translation. We select five different attributes from CelebAMask-HQ [54], including Gender, Mouth Slightly Open, Wavy Hair, No Beard, and Bald, as they cover human facial information from global structures to local details. The attribute translation task is formulated as a multi-domain UNIT task, where each domain corresponds to a positive/negative attribute. As shown in Fig. 7, our method generates vivid images and changes the attributes of input images accurately. It demonstrates that our method can work on UNIT tasks with some degree of perceptual structure change.

4.2.4 UNIT of mask and natural image

In Fig. 8, we show segmentation-to-RGB translation results of our method on six different ImageNet classes [59]. Thanks to the BigGAN-deep generator backbone [10], the synthesized images are highly realistic and distinct. The

synthesized images faithfully follow the shapes provided in segmentation masks. Although the input segmentation masks are imperfect, *e.g.*, the right ear of the cat is missing, our method learns a robust mapping between segmentation masks and RGB images to restore the unreasonable parts of masks.

4.3 Comparison with State-of-the-art Methods

Here, we compare the proposed method with the state-of-the-art methods on two UNIT settings: standard multi-domain UNIT (see Sec. 4.3.1) and domain-scalable UNIT (see Sec. 4.3.2). We also compare our method with three pretrained GANs-based UNIT methods in Sec. 4.3.3.

4.3.1 Standard multi-domain UNIT setting

Baseline methods. We compare five UNIT methods, including three representative dual-domain methods, *i.e.*, CycleGAN [1], CUT [3], and MUNIT [24], and two representative multi-domain methods, *i.e.*, StarGANv2 [8] and DRIT++ [6].




Fig. 9: A comparison of the state-of-the-art unpaired image translation methods on facial image domains. **(Top)** face-to-sketch results. **(Bottom)** segmentation-to-face results.



Fig. 10: A comparison of the unpaired image translation methods for translating segmentation masks to ImageNet images [59].

Quantitative results. Table 2 shows that our method achieves the best FID and LPIPS in most cases. CUT [3] shows the best performance for segmentation-to-face. We conjecture it is because the outputs of CUT [3] does not strictly match the input segmentation map and thus have a higher diversity. The results demonstrate the effectiveness of our method for standard UNIT tasks.

Qualitative results. Fig. 9 shows a qualitative comparison of the baseline methods for translating between multiple domains of facial images. The first row shows the results of translating an RGB image to a sketch (drawing). Both DRIT++ [6] and our method show reasonable results. DRIT++ is basically built upon feature style transfer, thus its synthesis images look sharper but contain more visual artifacts, as shown in the first example of Fig. 9. The second row shows the results of translating semantic segmentation to RGB images. Our method delivers the most vivid and realistic translation result. CUT [3] is the second-best method for segmentation to RGB images due to its effective contrastive learning framework.

In Fig. 10, we qualitatively compare our method with the dual-domain UNIT methods for translating foreground segmentation masks to ImageNet images. Our method delivers vivid dog and plane images. The baseline methods show poorer performances on this task, mainly because there are fewer training samples (*i.e.*, 1000 images per domain) in




Fig. 11: Results of the domain-scalable UNIT task, which aims at adding a new sketch drawing domain (as denoted in green bounding boxes) to existing UNIT models.

this experiment, compared against the segmentation-to-face experiment (27k images per domain).

4.3.2 Domain-scalable UNIT setting

The task is to extend an off-the-shelf UNIT model (*e.g.*, trained for segmentation maps and RGB images) to a new visual domain (*e.g.*, sketch drawings) provided a set of images of the new domain.

Baseline methods. We compare our method with a multi-domain method DRIT++ [6] and a unrestricted-domain method ComboGAN [35]. For the pretrained dual-domain DRIT++ [6] model, we fine-tune it with the data of all three available domains. For the pretrained dual-domain ComboGAN [35] model, we add an encoder and a decoder that optimize the cyclic consistency loss between the sketch domain and both two existing domains.

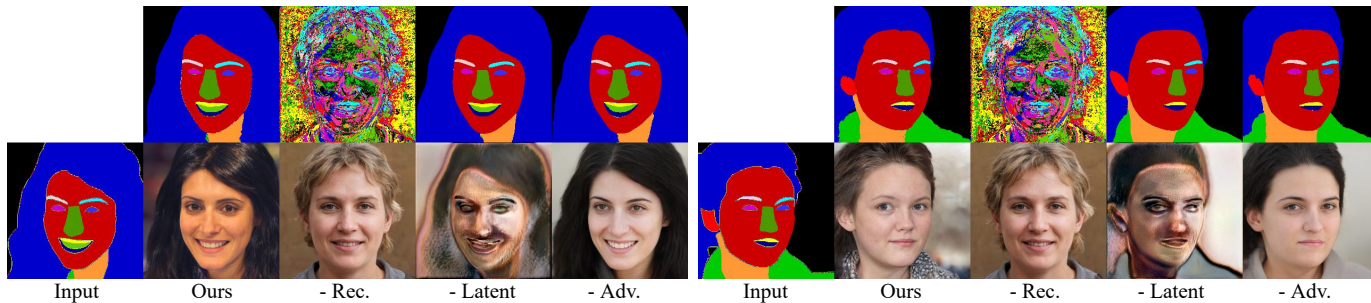


Fig. 12: An ablation study of our training objectives by removing one loss at a time. We show the **(top)** reconstructed segmentation masks and the **(bottom)** generated RGB images. “- Rec”: removing the reconstruction loss in Eq. 6. “- Latent”: removing the latent regularization loss in Eq. 3. “- Adv.”: removing the adversarial loss in Eq. 4.




Fig. 13: A comparison of pretrained GANs-based unpaired image translation methods for RGB-to-sketch and RGB-to-segmentation translations.

Qualitative results. Fig. 11 shows qualitative results of the domain-scalable UNIT setting. First, the proposed method generates more visually-pleasing facial images compared with baseline methods even on existing domains, *i.e.*, segmentation-to-RGB, thanks to the powerful pretrained GANs generator. Moreover, for the newly added sketch domain, sketch images translated by the proposed method are better aligned to the input segmentation masks than the baseline methods. The two baseline methods cannot well align the perceptual structure of the new domain to those of previous domains. The results validate the superior scalability of our method for new visual domains.

4.3.3 Pretrained GANs-based UNIT methods

Baseline methods. FreezeG [51] and Cartoon-StyleGAN [66] are pretrained GANs-based UNIT methods. To extend to new visual domains, they freeze the lower layers of GANs generators while fine-tuning the other layers. To make them be able to encode images into GAN’s latent space, we adopt an effective optimization-based GAN inversion method [67]. **Qualitative results.** Fig. 13 shows that the two baseline methods fail to achieve a good visual alignment of inputs and predictions. For the face-to-sketch translation, all training sketches are of frontal faces in the CUFSS dataset [55], such that the adversarial learning scheme adopted by the baseline methods limits them to generate sketches of frontal faces only. However, our method is based on feature reconstruction that the perceptual structures of generated images well follow those of the inputs. For face-to-segmentation UNIT, the hair segmentation maps pre-

dicted by the baseline methods are not accurate enough that the hairs are mistakenly recognized as shadows around faces. Our method predicts a more accurate segmentation map. It demonstrates the effectiveness of our latent space anchoring algorithm for aligning visual domains via the latent space of pretrained GANs.

4.4 Ablation Study

We further conduct various ablation experiments to examine the effects of various pipeline design choices.

4.4.1 Learning objectives

Fig. 12 studies the effects of the three loss terms in our training objectives, including the reconstruction loss, latent regularization loss, and adversarial loss, by separately removing one of them during training. Fig. 12 shows that the reconstruction loss is vital for our framework, as removing it results in a fail reconstruction of both images and segmentation masks. It validates that the reconstruction loss plays an key role in aligning the structures of images from different domains. Removing the latent regularization loss results in the distortion of generated images, because the encoder fails to embed input images into the latent space of GANs. The adversarial loss enhances the fidelity of the generated images, making them contain richer colors, textures, and background information. The results shown in Fig. 12 are in line with our intuition for devising the three learning objectives.

4.4.2 Training iteration

Fig. 14 visualizes the evolution of model results during training. Within the first 400 training iterations, the generated images are similar to that of iteration-0, *i.e.*, the initialized model. Beginning from iteration-400, the perceptual structures of reconstructed segmentation masks and generated images change to the input segmentation gradually. At iteration-3k, the reconstructed segmentation is basically identical to the input segmentation. Comparing iteration-3k and iteration-200k, the reconstructed segmentation masks are similar, while the image of iteration-200k shows richer appearance, color, and background details. It denotes that the model can learn latent space anchoring quickly, then improves latent space encoding afterwards.

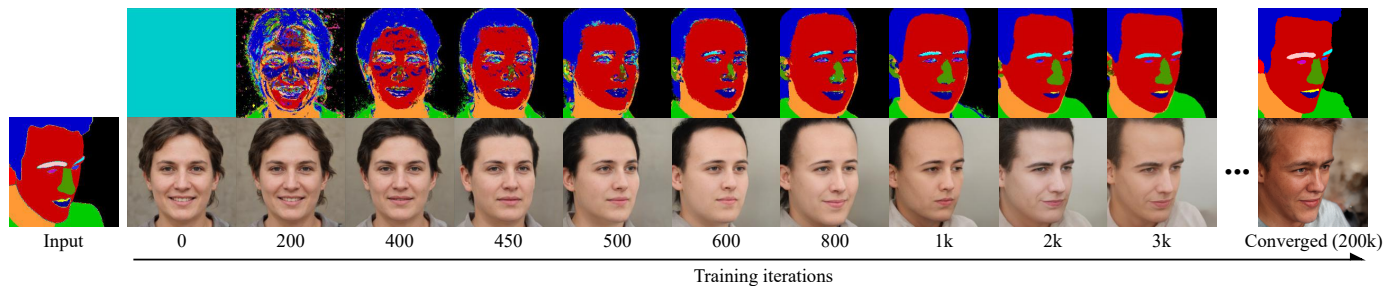


Fig. 14: The evolution of **(top)** reconstructed masks and **(bottom)** generated RGB images during training.

TABLE 3: Ablation studies of the input feature map size and network architecture for the regressor. Results are reported on segmentation-to-RGB task with a StyleGAN2 generator backbone [9] pretrained on FFHQ dataset [53]. The “All” in the column of Feature size denotes all the feature maps from 4×4 to 256×256 are concatenated as the input of regressor.

Method	Feature size	Regressor	FID↓	KID $\times 10^3$ ↓	SSIM↑	LPIPS↓
Latent Anchor (ours)	256×256	6 conv layers	83.8	72.8	0.307	0.442
Ablation of feature size	4×4	-	201.5	199.0	0.207	0.525
	16×16	-	188.9	217.7	0.269	0.476
	64×64	-	115.6	111.5	0.260	0.493
	128×128	-	116.2	116.2	0.267	0.481
	All	-	129.5	129.8	0.279	0.484
Ablation of regressor	-	3 conv layers	93.8	83.5	0.297	0.451
	-	6 ResBlocks	155.1	151.0	0.249	0.579



Fig. 15: A comparison of generator backbones pretrained on different datasets for unpaired image translation from *AFHQ-wild* to *AFHQ-cat* [8].

4.4.3 Pretrained GANs Prior

We study the effect of generator backbones pretrained on different datasets for the translation from *AFHQ-wild* to *AFHQ-cat* [8]. As shown in Fig. 15, different pretrained generator priors lead to different visual results, while all three GAN priors deliver decent image translation results, indicating the robustness of our method. In addition, when the GAN generator is pretrained on either the input or output domain, the results would contain slightly more fine-grained details. For example, the *AFHQ-cat* and *AFHQ-wild* results in Fig. 15 show more texture details on the cat faces.

4.4.4 Input feature maps of Regressor

We further study which feature maps of a StyleGAN2 [9] are more effective for the proposed method. Table 3 shows that the translation performances are better with higher-level feature maps, because the low-level feature maps contain fewer perceptual structures, leading to a worse latent space anchoring effect. Concatenating all the generator feature maps from 4×4 to 256×256 also achieves a poor performance. We conjecture it is because there are too many fea-

ture channels, which would lead to an over-strong regressor that fails the latent space anchoring process.

4.4.5 Network architecture of Regressor

As discussed above, the latent space anchoring method limits the representation ability of regressors. However, a very weak regressor would not well decode the feature maps to multi-domain images. The bottom of Table 3 studies different regressor network architectures, including 3 conv layers, 6 conv layers, and 6 ResNet blocks [62]. The 6 conv layers perform the best. 3 conv layers perform slightly worse than 6 conv layers. 6 ResBlocks are much more powerful than the aforementioned two networks, but they show a poor performance since they corrupt the weak regressor assumption.

4.5 User Study

We conduct a user study to subjectively compare the performance of the proposed method with the state-of-the-art methods. We choose the *segmentation map* \rightarrow *face* UNIT task, and randomly collect 300 sets of samples from the image translation results to form a sample pool. For each user, we randomly select 15 sets of samples from the pool, where each set consists of an input image and the translated images from six comparison methods (CycleGAN [1], CUT [3], MUNIT [24], DRIT++ [6], ComboGAN [35], and our method). For each set, we let the user choose one image of the best visual quality and one image that can best match the structure of the input. To ensure a fair user study, we design a website and place all the results of different algorithms one-by-one, where the order of images is shuffled individually for each set of results. We collected 24 questionnaires with 360 votes in total. Table 4 shows the result of our user study. Our method obtains 81.39% votes regarding the best visual quality and 39.72% votes regarding

TABLE 4: A user study on *segmentation map* \rightarrow *face* task. “Visual Quality Votes”: the percentage of a method that obtains the best visual quality votes. “Structure Match Votes”: the percentage of a method that obtains the best structure match votes.

Method	ComboGAN [35]	CUT [3]	CycleGAN [1]	DRIT++ [6]	MUNIT [24]	Latent Anchor (ours)
Visual Quality Votes (%)	1.67	10.56	0.56	4.72	1.11	81.39
Structure Match Votes (%)	4.44	27.22	4.44	16.39	7.78	39.72



Fig. 16: Generating high-resolution high-fidelity multi-domain and multimodal images via sampling from the latent space of StyleGAN2-FFHQ [9]. In each example, upper left shows the image generated by originally encoded W^+ . The rest three RGB images are the corresponding multimodal generations obtained by replacing the 5-th to 18-th (left) and 7-th to 18-th (right) W^+ of StyleGAN2 [9] with random variables.



Fig. 17: A progressive synthesis of human faces from coarse-grained to fine-grained semantic segmentation masks. Our method can translate the multi-domain inputs easily via anchoring them to the latent space of GANs.

the best structure match, both ranking the first among all the comparison methods. It demonstrates that our method can deliver images of good quality and preserve the structure of input images in terms of human evaluation.

4.6 Additional Applications

4.6.1 Multimodal multi-domain dataset synthesis

Since our method aligns multi-domain images to the latent space of pretrained powerful GANs, we can generate infinite well-aligned, high-resolution, and highly realistic multimodal multi-domain images by sampling from the latent space of generators. As shown in Fig. 16, the landmark maps, semantic segmentation maps, sketch drawings are generated along with the RGB images (the upper left ones). In addition to multi-domain generation, our method also supports generating *multimodal* images (*i.e.*, images with different appearance details as presented in MUNIT [24] and DRIT [25]) by replacing the W^+ latent variables of StyleGAN2 [9] with randomly sampled variables. Fig. 16 shows the semantic structures of multi-modal RGB images are highly aligned with the multi-domain generations. The multi-domain and multimodal image generation would benefit many downstream visual analysis tasks, such as

generating a large number of data samples to augment existing multi-domain image datasets.

4.6.2 Progressive image translation

Different users may prefer using inputs of various annotation levels for image translation. Here we show a progressive image translation example in Fig. 17, where the segmentation masks vary from coarse-grained to fine-grained annotations. We train encoders and regressors for these four types of segmentation inputs individually. Thanks to latent space anchoring, we unify multiple image domains easily by anchoring new visual domains to the latent space of GANs. In inference, more fine-grained segmentation masks are inferred based on previous ones, providing a series of segmentation maps of different annotation levels to users.

5 CONCLUSION

In this paper, we address a new challenge of unpaired image translation, *i.e.*, domain-scalable UNIT, and presented a novel latent space anchoring method to address this new challenge. Our method aligns different visual domains to the latent space of pretrained GANs generator via learning to reconstruct images on single domains. Intuitive analysis has interpreted the proposed method by connecting the

minimization of reconstruction error to the minimization of semantic structure gap between inputs and translation results. We have comprehensively evaluated our method on diverse domains including human faces, animal faces, and ImageNet images with StyleGAN2 [9] and BigGAN-deep [10] as the generator backbones. Quantitative, qualitative, and user study results have validated that our method achieves superior performance compared with previous UNIT methods for standard and domain-scalable UNIT tasks.

6 LIMITATION AND ETHICAL STATEMENT

Limitation. One limitation of our method is that it is devised for aligning perceptual structures of images. Therefore, only visual domains having similar perceptual structures can be incorporated into this framework. For instance, human facial images, human facial segmentation masks, and animal facial images can be aligned. However, human facial images and ImageNet foreground segmentation masks would not be well aligned as they have different perceptual structures in datasets. In the future, we plan to resolve this limitation by incorporating deformable models [68], [69], semantic layouts [70], [71], or geometric priors [72], [73].


Ethical Statement. The proposed method is a general multi-domain image translation framework. As shown by the experimental results, our method can be used in human and animal face translation, which is considered a potential ethical issue. To prevent the abusive usage of the proposed method, the codes and model checkpoints of the human/animation translation are not released to the public. They are only released to researchers upon formal requests and can only be used for research purposes.

REFERENCES

- [1] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *ICCV*, 2017, pp. 2223–2232.
- [2] M.-Y. Liu, T. Breuel, and J. Kautz, "Unsupervised image-to-image translation networks," *NeurIPS*, vol. 30, 2017.
- [3] T. Park, A. A. Efros, R. Zhang, and J.-Y. Zhu, "Contrastive learning for unpaired image-to-image translation," in *ECCV*, 2020, pp. 319–345.
- [4] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *CVPR*, 2017, pp. 1125–1134.
- [5] T. Park, M.-Y. Liu, T.-C. Wang, and J.-Y. Zhu, "Semantic image synthesis with spatially-adaptive normalization," in *ICCV*, 2019, pp. 2337–2346.
- [6] H.-Y. Lee, H.-Y. Tseng, Q. Mao, J.-B. Huang, Y.-D. Lu, M. Singh, and M.-H. Yang, "Drit++: Diverse image-to-image translation via disentangled representations," *IJCV*, vol. 128, no. 10, pp. 2402–2417, 2020.
- [7] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo, "Stargan: Unified generative adversarial networks for multi-domain image-to-image translation," in *CVPR*, 2018, pp. 8789–8797.
- [8] Y. Choi, Y. Uh, J. Yoo, and J.-W. Ha, "Stargan v2: Diverse image synthesis for multiple domains," in *CVPR*, 2020, pp. 8188–8197.
- [9] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila, "Analyzing and improving the image quality of stylegan," in *CVPR*, 2020, pp. 8110–8119.
- [10] A. Brock, J. Donahue, and K. Simonyan, "Large scale gan training for high fidelity natural image synthesis," in *ICLR*, 2018.
- [11] A. Almahairi, S. Rajeshwar, A. Sordoni, P. Bachman, and A. Courville, "Augmented cyclegan: Learning many-to-many mappings from unpaired data," in *ICML*, 2018, pp. 195–204.

- [12] J. Kim, M. Kim, H. Kang, and K. Lee, "U-gat-it: Unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation," *arXiv preprint arXiv:1907.10830*, 2019.
- [13] Y. Zhao, R. Wu, and H. Dong, "Unpaired image-to-image translation using adversarial consistency loss," in *ECCV*, 2020, pp. 800–815.
- [14] R. Chen, W. Huang, B. Huang, F. Sun, and B. Fang, "Reusing discriminators for encoding: Towards unsupervised image-to-image translation," in *CVPR*, 2020, pp. 8168–8177.
- [15] O. Nizan and A. Tal, "Breaking the cycle-colleagues are all you need," in *CVPR*, 2020, pp. 7860–7869.
- [16] X. Shao and W. Zhang, "Spatchgan: A statistical feature based discriminator for unsupervised image-to-image translation," in *ICCV*, 2021, pp. 6546–6555.
- [17] Y. Zhao and C. Chen, "Unpaired image-to-image translation via latent energy transport," in *CVPR*, 2021, pp. 16418–16427.
- [18] F. Pizzati, P. Cerri, and R. de Charette, "Comogan: continuous model-guided image-to-image translation," in *CVPR*, 2021, pp. 14288–14298.
- [19] T. Kim, M. Cha, H. Kim, J. K. Lee, and J. Kim, "Learning to discover cross-domain relations with generative adversarial networks," in *ICML*, 2017, pp. 1857–1865.
- [20] Z. Yi, H. Zhang, P. Tan, and M. Gong, "Dualgan: Unsupervised dual learning for image-to-image translation," in *ICCV*, 2017, pp. 2849–2857.
- [21] Y. Liu, M. De Nadai, J. Yao, N. Sebe, B. Lepri, and X. Alameda-Pineda, "Gmm-unit: Unsupervised multi-domain and multi-modal image-to-image translation via attribute gaussian mixture modeling," *arXiv preprint arXiv:2003.06788*, 2020.
- [22] W. Xu and G. Wang, "A domain gap aware generative adversarial network for multi-domain image translation," *IEEE TIP*, vol. 31, pp. 72–84, 2021.
- [23] V. Vinod, K. R. Prabhakar, R. V. Babu, and A. Chakraborty, "Multi-domain conditional image translation: Translating driving datasets from clear-weather to adverse conditions," in *ICCV*, 2021, pp. 1571–1582.
- [24] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz, "Multimodal unsupervised image-to-image translation," in *ECCV*, 2018, pp. 172–189.
- [25] H.-Y. Lee, H.-Y. Tseng, J.-B. Huang, M. Singh, and M.-H. Yang, "Diverse image-to-image translation via disentangled representations," in *ECCV*, 2018, pp. 35–51.
- [26] B. Zhao, B. Chang, Z. Jie, and L. Sigal, "Modular generative adversarial networks," in *ECCV*, 2018, pp. 150–165.
- [27] A. Pumarola, A. Agudo, A. M. Martinez, A. Sanfeliu, and F. Moreno-Noguer, "Ganimation: Anatomically-aware facial animation from a single image," in *ECCV*, 2018, pp. 818–833.
- [28] M.-Y. Liu, X. Huang, A. Mallya, T. Karras, T. Aila, J. Lehtinen, and J. Kautz, "Few-shot unsupervised image-to-image translation," in *ICCV*, 2019, pp. 10551–10560.
- [29] A. Romero, P. Arbeláez, L. Van Gool, and R. Timofte, "Smit: Stochastic multi-label image-to-image translation," in *ICCV Workshops*, 2019, pp. 0–0.
- [30] Y. Wang, A. Gonzalez-Garcia, J. van de Weijer, and L. Herranz, "Sdit: Scalable and diverse cross-domain image translation," in *Proceedings of the 27th ACM International Conference on Multimedia*, 2019, pp. 1267–1276.
- [31] Y.-C. Chen, X. Xu, and J. Jia, "Domain adaptive image-to-image translation," in *CVPR*, 2020, pp. 5274–5283.
- [32] K. Baek, Y. Choi, Y. Uh, J. Yoo, and H. Shim, "Rethinking the truly unsupervised image-to-image translation," in *ICCV*, 2021, pp. 14154–14163.
- [33] Y. Liu, E. Sangineto, Y. Chen, L. Bao, H. Zhang, N. Sebe, B. Lepri, W. Wang, and M. De Nadai, "Smoothing the disentangled latent style space for unsupervised image-to-image translation," in *CVPR*, 2021, pp. 10785–10794.
- [34] Y. Pang, J. Lin, T. Qin, and Z. Chen, "Image-to-image translation: Methods and applications," *IEEE Transactions on Multimedia*, 2021.
- [35] A. Anoosheh, E. Agustsson, R. Timofte, and L. Van Gool, "Comogan: Unrestrained scalability for image domain translation," in *CVPR Workshops*, 2018, pp. 783–790.
- [36] D. Bau, J.-Y. Zhu, H. Strobel, B. Zhou, J. B. Tenenbaum, W. T. Freeman, and A. Torralba, "Gan dissection: Visualizing and understanding generative adversarial networks," in *ICLR*, 2019.
- [37] W. Xia, Y. Zhang, Y. Yang, J.-H. Xue, B. Zhou, and M.-H. Yang, "Gan inversion: A survey," *arXiv preprint arXiv:2101.05278*, 2021.

- [38] R. Abdal, Y. Qin, and P. Wonka, "Image2stylegan: How to embed images into the stylegan latent space?" in *ICCV*, 2019, pp. 4432–4441.
- [39] E. Richardson, Y. Alaluf, O. Patashnik, Y. Nitzan, Y. Azar, S. Shapiro, and D. Cohen-Or, "Encoding in style: a stylegan encoder for image-to-image translation," in *CVPR*, 2021, pp. 2287–2296.
- [40] Y. Shen, J. Gu, X. Tang, and B. Zhou, "Interpreting the latent space of gans for semantic face editing," in *CVPR*, 2020, pp. 9243–9252.
- [41] A. Cherepkov, A. Voynov, and A. Babenko, "Navigating the gan parameter space for semantic image editing," in *CVPR*, 2021, pp. 3671–3680.
- [42] O. Tov, Y. Alaluf, Y. Nitzan, O. Patashnik, and D. Cohen-Or, "Designing an encoder for stylegan image manipulation," *ACM TOG*, vol. 40, no. 4, pp. 1–14, 2021.
- [43] X. Pan, X. Zhan, B. Dai, D. Lin, C. C. Loy, and P. Luo, "Exploiting deep generative prior for versatile image restoration and manipulation," *IEEE TPAMI*, 2021.
- [44] R. Abdal, Y. Qin, and P. Wonka, "Image2stylegan++: How to edit the embedded images?" in *CVPR*, 2020, pp. 8296–8305.
- [45] Y. Shen and B. Zhou, "Closed-form factorization of latent semantics in gans," in *CVPR*, 2021, pp. 1532–1540.
- [46] A. Voynov and A. Babenko, "Unsupervised discovery of interpretable directions in the gan latent space," in *ICML*, 2020, pp. 9786–9796.
- [47] Y. Zhang, H. Ling, J. Gao, K. Yin, J.-F. Lafleche, A. Barriuso, A. Torralba, and S. Fidler, "Datasetgan: Efficient labeled data factory with minimal human effort," in *CVPR*, 2021, pp. 10145–10155.
- [48] D. Pakhomov, S. Hira, N. Wagle, K. E. Green, and N. Navab, "Segmentation in style: Unsupervised semantic image segmentation with stylegan and clip," *arXiv preprint arXiv:2107.12518*, 2021.
- [49] S. Benaim and L. Wolf, "One-shot unsupervised cross domain translation," *advances in neural information processing systems*, vol. 31, 2018.
- [50] J. Back, "Fine-tuning stylegan2 for cartoon face generation," *arXiv preprint arXiv:2106.12445*, 2021.
- [51] S. Mo, M. Cho, and J. Shin, "Freeze the discriminator: a simple baseline for fine-tuning gans," *arXiv preprint arXiv:2002.10964*, 2020.
- [52] K. et al., "Alias-free generative adversarial networks," in *NeurIPS*, 2021.
- [53] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *CVPR*, 2019, pp. 4401–4410.
- [54] C.-H. Lee, Z. Liu, L. Wu, and P. Luo, "Maskgan: Towards diverse and interactive facial image manipulation," in *CVPR*, 2020, pp. 5549–5558.
- [55] X. Wang and X. Tang, "Face photo-sketch synthesis and recognition," *IEEE TPAMI*, vol. 31, no. 11, pp. 1955–1967, 2008.
- [56] W. Zhang, X. Wang, and X. Tang, "Coupled information-theoretic encoding for face photo-sketch recognition," in *CVPR*, 2011, pp. 513–520.
- [57] V. Kazemi and J. Sullivan, "One millisecond face alignment with an ensemble of regression trees," in *CVPR*, 2014, pp. 1867–1874.
- [58] D. E. King, "Dlib-ml: A machine learning toolkit," *JMLR*, vol. 10, pp. 1755–1758, 2009.
- [59] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *CVPR*, 2009, pp. 248–255.
- [60] D. Kuettel, M. Guillaumin, and V. Ferrari, "Segmentation propagation in imagenet," in *ECCV*, 2012, pp. 459–473.
- [61] M. Guillaumin, D. Küttel, and V. Ferrari, "Imagenet auto-annotation with segmentation propagation," *IJCV*, vol. 110, no. 3, pp. 328–348, 2014.
- [62] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016, pp. 770–778.
- [63] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR*, 2015.
- [64] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "Gans trained by a two time-scale update rule converge to a local nash equilibrium," *NeurIPS*, vol. 30, 2017.
- [65] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *CVPR*, 2018, pp. 586–595.
- [66] J. Back, "Fine-tuning stylegan2 for cartoon face generation," *arXiv preprint arXiv:2106.12445*, 2021.
- [67] J. Zhu, Y. Shen, D. Zhao, and B. Zhou, "In-domain gan inversion for real image editing," in *ECCV*, 2020.
- [68] M. Jaderberg, K. Simonyan, A. Zisserman *et al.*, "Spatial transformer networks," *NeurIPS*, vol. 28, 2015.
- [69] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei, "Deformable convolutional networks," in *ICCV*, 2017, pp. 764–773.
- [70] H. Zheng, Z. Lin, J. Lu, S. Cohen, J. Zhang, N. Xu, and J. Luo, "Semantic layout manipulation with high-resolution sparse attention," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [71] H. Zheng, Z. Lin, J. Lu, S. Cohen, E. Shechtman, C. Barnes, J. Zhang, N. Xu, S. Amirghodsi, and J. Luo, "Cm-gan: Image inpainting with cascaded modulation gan and object-aware training," 2022.
- [72] Y. Tan, H. Zheng, Y. Zhu, X. Yuan, X. Lin, D. Brady, and L. Fang, "Crossnet++: Cross-scale large-parallax warping for reference-based super-resolution," *IEEE TPAMI*, vol. 43, no. 12, pp. 4291–4305, 2020.
- [73] P. Zhang, B. Zhang, D. Chen, L. Yuan, and F. Wen, "Cross-domain correspondence learning for exemplar-based image translation," in *CVPR*, 2020, pp. 5143–5153.



Siyu Huang received the B.E. degree and Ph.D. degree in information and communication engineering from Zhejiang University, Hangzhou, China, in 2014 and 2019. He is currently a Post-doctoral Fellow in the John A. Paulson School of Engineering and Applied Sciences at Harvard University. Before that, he was a Visiting Scholar at Language Technologies Institute in the School of Computer Science, Carnegie Mellon University in 2018, a Research Scientist at Big Data Laboratory, Baidu Research from 2019 to 2021, and a Research Fellow in the School of Electrical and Electronic Engineering at Nanyang Technological University in 2021. He has published more than 20 papers on top-tier computer science journals and conferences. His research interests are primarily in computer vision, multimedia analysis, and generative models.



Jie An received his Bachelor of Science (B.Sc) in Information and Computing Sciences and Master of Science (M.Sc) in Applied Mathematics, both from Peking University. He is currently a PhD student in Computer Science at the University of Rochester. His research interests include computer Vision, generative models, and AI+Art.



Donglai Wei received his Bachelor of Science (B.Sc) in Mathematics from Brown University and PhD degree in Computer Science from MIT. He is an assistant professor of computer science at Boston College. Before joining Boston College, he worked as a postdoctoral researcher at Harvard University. His research interests include brain image analysis and video understanding.



Zudi Lin received his PhD degree in Computer Science in 2022 from the John A. Paulson School of Engineering and Applied Sciences at Harvard University. Before that, he obtained M.S. in Computer Science from Harvard University in 2020 and B.S. in Biological Science from Tsinghua University in 2017. His research interests include deep learning, computer vision, and neuroscience.



Jiebo Luo (Fellow, IEEE) is Professor of Computer Science at the University of Rochester which he joined in 2011 after a prolific career of fifteen years at Kodak Research Laboratories. He has authored well over 500 technical papers and holds over 90 U.S. patents. His research interests include computer vision, NLP, machine learning, data mining, computational social science, and digital health. He has been involved in numerous technical conferences, including serving as program co-chair of ACM Multimedia 2010, IEEE CVPR 2012, ACM ICMR 2016, and IEEE ICIP 2017, as well as general co-chair of ACM Multimedia 2018. He has served on the editorial boards of the IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), IEEE Transactions on Multimedia (TMM), IEEE Transactions on Circuits and Systems for Video Technology (TCSVT), IEEE Transactions on Big Data (TBD), ACM Transactions on Intelligent Systems and Technology (TIST), Pattern Recognition, Knowledge and Information Systems (KAIS), Machine Vision and Applications, and Intelligent Medicine. He is the current Editor-in-Chief of the IEEE Transactions on Multimedia. Professor Luo is also a Fellow of NAI, ACM, AAAI, SPIE, and IAPR.



Hanspeter Pfister (Fellow, IEEE) received the MS degree in electrical engineering from ETH Zurich, Switzerland, and the PhD degree in computer science from the State University of New York at Stony Brook. He is currently a professor of computer science at the Harvard John A. Paulson School of Engineering and Applied Sciences and an affiliate faculty member of the Center for Brain Science. His research in visual computing lies at the intersection of visualization, computer graphics, and computer vision and spans a wide range of topics, including biomedical image analysis and visualization, image and video analysis, interpretable machine learning, and visual analytics in data science. From 2013 to 2017, he was the director of the Institute for Applied Computational Science. Before joining Harvard, he worked for over a decade at Mitsubishi Electric Research Laboratories, where he was an associate director and senior research scientist. He was the chief architect of VolumePro, Mitsubishi Electric's award-winning real-time volume rendering graphics card, for which he received the Mitsubishi Electric President's Award, in 2000. He was elected as an ACM fellow, in 2019, and an IEEE fellow, in 2022.