UniHDA: A Unified and Versatile Framework for Multi-Modal Hybrid Domain Adaptation

Hengjia Li[1], Yang Liu[1], Yuqi Lin[1], Zhanwei Zhang[1], Yibo Zhao[1], Weihang Pan[1], Tu Zheng[2], Zheng Yang[2], Chunjiang Yu[3], Boxi Wu[1], and Deng Cai[1]

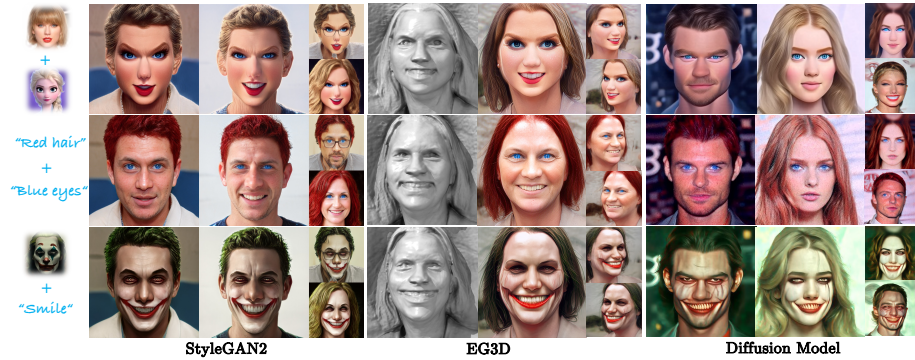[1] Zhejiang University   [2] Fabu Inc.   [3] Ningbo Port

Fig. 1: Given a pre-trained source generator and multiple target domains, UniHDA adapts the generator to a hybrid target domain that blends all characteristics at once and maintains robust cross-domain consistency. UniHDA supports both image and text modalities and is versatile to multiple generators.

Abstract. Recently, generative domain adaptation has achieved remarkable progress, enabling us to adapt a pre-trained generator to a new target domain. However, existing methods simply adapt the generator to a single target domain and are limited to a single modality, either text-driven or image-driven. Moreover, they cannot maintain well consistency with the source domain, which impedes the inheritance of the diversity. In this paper, we propose UniHDA, a **unified** and **versatile** framework for generative hybrid domain adaptation with multi-modal references from multiple domains. We use CLIP encoder to project multi-modal references into a unified embedding space and then linearly interpolate the direction vectors from multiple target domains to achieve hybrid domain adaptation. To ensure **consistency** with the source domain, we propose a novel cross-domain spatial structure (CSS) loss that maintains detailed spatial structure information between source and target generator. Experiments show that the adapted generator can synthesise realistic images with various attribute compositions. Additionally, our framework is generator-agnostic and versatile to multiple generators, *e.g.*, StyleGAN, EG3D, and Diffusion Models.

Keywords: Generative Domain Adaptation · Multi-Modal Adaptation · Hybrid Domain Adaptation · Generative Models

Fig. 2: Existing methods like NADA [8] fail to maintain consistency with the source domain for hybrid domain adaptation, resulting in overfitting to the limited references and impeding the inheritance of the diversity in the source domain.

1   Introduction

Benefiting from tremendous success of modern image generator [3, 12, 35, 43], generative domain adaptation has achieved remarkable progress during the past few years. Typically, it aims to adapt a pre-trained generator to a new target domain while preserving the variation in the source domain, *e.g.*, from the *human* domain to the *baby* domain. Depending on the modality of references, generative domain adaptation can be categorized into two schools: text-driven [2, 8, 14, 15, 19, 23, 24, 27, 40, 56] and image-driven [1, 7, 17, 22, 25, 26, 29, 45, 47, 49, 51–55].

Despite their promising results, there are still some limitations. A major limitation is that they only support adaptation from a source domain to individual target domain. These methods fail to directly adapt a *human* generator to more practical real-world scenarios like *person with red hair and blue eyes* or a combination of *Taylor* and *Elsa* (Fig. 1). For more general purposes, previous methods also fail with multi-modal references, *e.g.*, *smiling Joker* given the *non-smiling Joker* image and the *smile* text.

Another limitation of existing approaches is that, they are prone to overfit domain-specific attributes, especially in hybrid domain adaptation (Fig. 2). The main reason is that they do not maintain well consistency with the source domain, *e.g.*, posture and identity. This results in overfitting to the limited references and losing the diversity in the source domain.

In this paper, we propose UniHDA, a **Uni**fied and versatile framework for generative **H**ybrid **D**omain **A**daptation with multi-modal references from multiple domains. UniHDA facilitates the references of individual image and text prompt simultaneously and blends the attributes from target domains to create a hybrid domain. To enable multiple modalities, we leverage pre-trained CLIP [34] to project multi-modal references into a unified embedding space and represent the domain shift by the direction vector from the source embedding to the target embeddings.
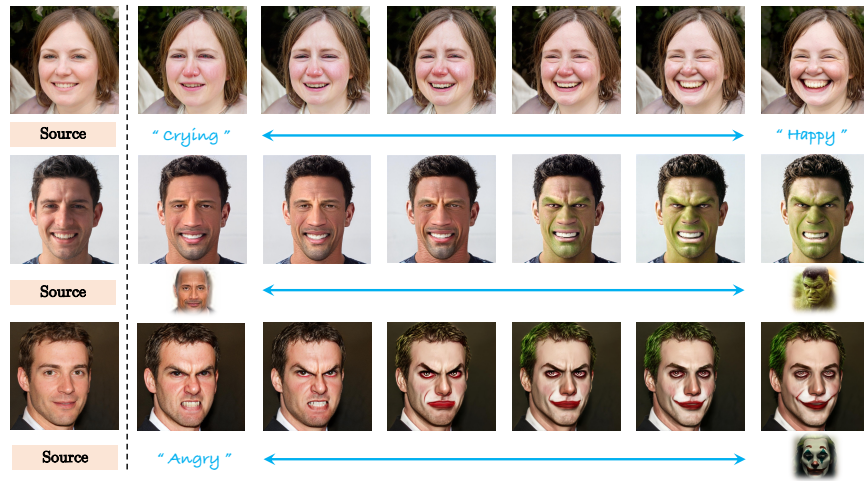
Fig. 3: Linear interpolation between multi-modal direction vectors. We represent the domain shift by the direction vector from source embedding to the target (*e.g.*, *Crying* or *Happy*). Linear interpolation of them during training will result in a smooth traversal. The coefficients for the right domain are respectively 0, 0.2, 0.4, 0.6, 0.8, and 1, while for the left domain, they are set inversely.

To achieve hybrid domain adaptation, we draw inspiration from the compositional capabilities in the latent space of StyleGAN [9, 38, 48]. We demonstrate that a semantically meaningful linear interpolation between direction vectors in CLIP's embedding space can uncover favorable compositional capabilities (Fig. 3). In light of this intriguing finding, we linearly interpolate direction vectors of multiple target domains to obtain the direction vector corresponding to the hybrid domain that semantically integrates attributes from all target domains.

Furthermore, we introduce a novel cross-domain spatial structure loss (CSS) to preserve the consistency between source and target generator by maintaining detailed spatial structure information. Concretely, we leverage pre-trained Dino-ViT [6, 31] to encode generated images into patch tokens with fine-grained spatial information. For cross-domain consistency, we maintain the correspondence between source and target tokens with contrastive learning [30].

We conduct experiments for a wide range of source and target domains to validate the effectiveness of our method. Results demonstrate that the adapted generator can synthesise realistic images with various attribute compositions. In addition, we show that UniHDA is agnostic to the type of generators, *e.g.*, StyleGAN [11–13], EG3D [4], and Diffusion models [10, 16]. Our contributions are as follows:

– We propose a **unified** and **versatile** framework for generative domain adaptation, which enables multi-modal references and hybrid target domain, *e.g.*, text-text, image-image, and image-text. We demonstrate successful adaptation to diverse domains for various generators and illustrate its advantage over other methods.

- We demonstrate strong compositional capabilities of direction vectors in CLIP's embedding space. Taking advantage of it, we propose to linearly interpolate the direction vectors for multi-modal hybrid domain adaptation.
- We propose a cross-domain spatial structure loss to maintain consistency with source domain. It is conducted in generator-agnostic embedding space which is versatile for various generators, *e.g.*, StyleGAN, EG3D, and Diffusion models. To our knowledge, it is the very first trial in generative domain adaptation.

2   Related Work

Text-driven Generative Domain Adaptation. Text-driven domain adaptation [2, 8, 14, 15, 19, 23, 24, 27, 40, 56] involves using a textual prompt to shift the domain of a pre-trained model toward a new domain. For example, Style-NADA [8] presents a local direction CLIP [34] loss to align the embeddings of the generated images and text. Based on Style-NADA, Domain Expansion (DE) [28] proposes to expand the generator to jointly model multiple domains with texts.
Image-driven Generative Domain Adaptation. Image-driven generative domain adaptation [1, 7, 17, 22, 25, 26, 29, 45, 47, 49, 51–55] refers to the adaptation of a pre-trained image generator to a new target domain using a limited number of training images. Due to the scarcity of training images, prior researches often integrate additional regularization terms to prevent overfitting. For instance, CDC [29] introduces the instance distance consistency loss to maintain the distance between different instances in the source domain. DiFa [51] utilizes GAN inversion [42] to align the latent codes which helps inherit diversity from the source generator. Although these works have made significant strides in generative domain adaptation, they heavily rely on the discriminator or generator, making it challenging to handle hybrid domain adaptation and extend to other generators.
Generative Hybrid Domain Adaptation. To achieve hybrid domain adaptation, several domain adaptation methods propose to train a separate generative model per domain and combine their effects in test-time, *e.g.*, Style-NADA [8]. Although having potential for hybrid domain adaptation, it doubles the model size and training time due to the requirement for training multiple models separately. Domain Expansion (DE) [28] proposes to expand the generator to jointly model multiple domains via decompose latent space. However, it requires the source dataset (*e.g.*, FFHQ [12]) for regularization, which significantly increases training time. Furthermore, it relies on the semantic latent space of the generator (*e.g.*, StyleGAN [12] and DiffAE [33]) for hybrid domain adaptation, limiting its applicability to a broader range of generators. Recently, FHDA [20] proposes few-shot hybrid domain adaptation and introduces a directional subspace loss to achieve it. Differently, we focus on multi-modal references with the text and one-shot image, which offers greater flexibility and versatility.
Disentanglement in Generative Models.   As observed in StyleGAN [12], the latent space is essentially a linear subspace. Recent works [9, 32, 37, 38, 41, 44, 46, 48] propose to find individual latent factors for image variations. Among them, SeFa [38] computes the eigenvalues of the transformation matrix to find the latent

directions. As for diffusion models, DiffAE [33] explores the possibility of using DPMs for representation learning and seeks to extract a meaningful and decodable representation of an input image via autoencoding. Drawing inspiration from the disentanglement in the latent space of them, we demonstrate that a semantically meaningful linear interpolation between the direction vectors in CLIP's embedding space can similarly uncover favorable compositional capabilities, which facilitates us to achieve hybrid domain adaptation.

3   Method

3.1   Multi-Modal Hybrid Domain Adaptation

We start with a pre-trained generator $G_{\mathcal{S}}$ (*e.g.*, StyleGAN [11–13] and Diffusion model [10,39])), that maps from noise z to images in a source domain $\mathcal{S}$. Given a new target domain $\mathcal{T}$ referenced by texts [8,16,18,51] or images [22,25,26,29,47,53], generative domain adaptation aims to adapt $G_{\mathcal{S}}$ to yield a target generator $G_{\mathcal{T}}$, which can generate images similar to domain $\mathcal{T}$.

Despite the promising results of existing methods, a major limitation of them is that they only support adaptation from the source domain to individual target domains and fail to directly adapt the generator to the hybrid domain which blends the characteristics of multiple domains. Furthermore, they fail with multi-modal adaptation driven by texts and images simultaneously.

For more general purposes, we propose multi-modal hybrid domain adaptation. Given N domains $\{\mathcal{T}_i\}_{i=1}^{N}$ with one-shot image $\{Y_i\}$ and M domains $\{\mathcal{T}_j\}_{j=1}^{M}$ with the text prompt $\{P_j\}$, it aims to adapt the source generator $G_{\mathcal{S}}$ to $G_{\mathcal{T}}$ that models the hybrid domain $\mathcal{T} = \{\mathcal{T}_i\} \cup \{\mathcal{T}_j\}$ and generates images with integrated characteristics. To the end, we introduce UniHDA, a unified and versatile framework for multi-modal hybrid domain adaptation (Fig. 4).

3.2   Multi-modal Direction Loss

To enable multiple modalities, we leverage pre-trained CLIP [34] to encode text-image references into a unified semantic embedding space. Drawing inspiration from CLIP-based methods [8, 16, 18, 51], we represent the *domain shift* as the direction vector Δf_{dom} from the source embedding to the target embedding. For image reference Y_i and its CLIP embedding f_i, the *domain shift* is calculated by

$$\Delta f_{dom} = f_i - \overline{f_s}, \tag{1}$$

where $\overline{f_s}$ is the mean embedding of several samples generated by $G_{\mathcal{S}}$. For text prompt P_j and its CLIP embedding f_j,

$$\Delta f_{dom} = f_j - \widetilde{f_s}, \tag{2}$$

where $\widetilde{f_s}$ is the embedding of the source text prompt.

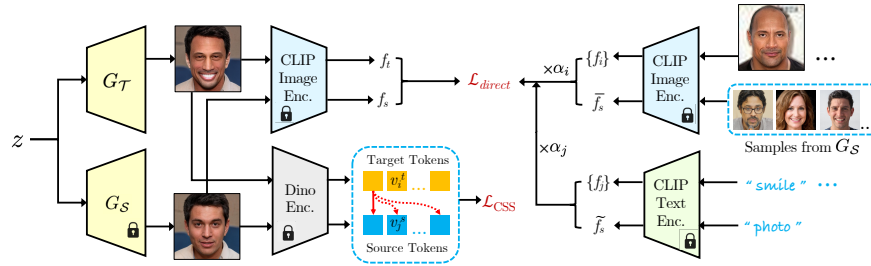**Fig. 4:** Overview of UniHDA with multi-modal direction loss $\mathcal{L}_{direct}$ and cross-domain spatial structure loss $\mathcal{L}_{\text{CSS}}$. Utilizing CLIP image encoder and text encoder, $\mathcal{L}_{direct}$ encourages $G_{\mathcal{T}}$ to faithfully acquire domain-specific characteristics with multi-modal references. To facilitate diversity inherited from $G_{\mathcal{S}}$, $\mathcal{L}_{\text{CSS}}$ improves cross-domain consistency by maintaining detailed spatial structure information. The red solid line represents positive pairs, while red dashed lines represent negative pairs.

To adapt $G_{\mathcal{S}}$, we initialize a new generator $G_{\mathcal{T}}$ from $G_{\mathcal{S}}$ and finetune it by aligning the *sample-shift* direction Δf_{samp} with the *domain-shift* direction Δf_{dom}. Formally,

$$\Delta f_{samp} = f_t - f_s,$$
$$\mathcal{L}_{direct} = 1 - \frac{\Delta f_{samp} \cdot \Delta f_{dom}}{\|\Delta f_{samp}\| \|\Delta f_{dom}\|}, \tag{3}$$

where f_s and f_t are the embeddings of samples generated by $G_{\mathcal{S}}$ and $G_{\mathcal{T}}$ with the same noise.

3.3 Linear Composition of Direction Vectors

To achieve the hybrid domain adaptation, we draw inspiration from the compositional capabilities in the latent space of StyleGAN [9,38,48]. We illustrate that a linear interpolation between two direction vectors in the embedding space of CLIP, which is semantically meaningful, reveals promising compositional capabilities. As shown in Fig. 3, we can smoothly interpolate between two direction vectors calculated by distinct target prompts and source prompt "*photo*", resulting in a gradual adaptation toward the target domain.

In light of this intriguing finding, we employ linear interpolation on the direction vectors of multi-modal target domains, to derive a consolidated direction vector representing the hybrid domain that semantically integrates all attributes. For given domain coefficients $\{\alpha_i\}$ and $\{\alpha_j\}$, we obtain the direction vector

$$\Delta f_{dom} = \sum_{i=1}^{N} \alpha_i (f_i - \overline{f_s}) + \sum_{j=1}^{M} \alpha_j (f_j - \widetilde{f_s}), \tag{4}$$

which represents the *domain shift* between the hybrid domain and source domain. We then substitute Eq. (4) into Eq. (3) to adapt $G_{\mathcal{S}}$ to the hybrid domain.

3.4 Cross-domain Spatial Structure Loss

Albeit the direction loss achieves multi-modal hybrid domain adaptation, the adapted generator are prone to overfit domain-specific attributes. This exacerbates when it comes to image-image and image-text scenarios owing to the scarcity of the images. To address this issue, we introduce a novel cross-domain spatial structure loss (CSS) to enhance cross-domain consistency, ensuring the preservation of intricate spatial structural information between the source and target generator.

Specifically, we leverage the pre-trained Dino-ViT [6, 31] to encode the generated images into patch tokens, containing detailed spatial structural information. Dino-ViT is self-supervised trained to focus on the distinction between subjects of the same class [36], which facilitates us to maintain the cross-domain consistency. Motivated by contrastive learning [30], we reduce the distance between the positive token pairs at the same position and push away the negative token pairs at the other positions by

$$\mathcal{L}_{\mathrm{CSS}} = -\sum_i \log \frac{\exp(v_i^t \cdot v_i^s)}{\sum_j \exp(v_i^t \cdot v_j^s)}, \tag{5}$$

where v_i^t and v_j^s are the i-th and j-th tokens in the last layer of Dino-ViT from $G_{\mathcal{T}}$ and $G_{\mathcal{S}}$ respectively. The dot mark $\cdot$ represents dot product.

Overall, our training loss consists of two terms, i.e., the multi-modal direction loss $\mathcal{L}_{direct}$ to achieve multi-modal hybrid domain adaptation and the cross-domain spatial structure loss $\mathcal{L}_{\mathrm{CSS}}$ to maintain cross-domain consistency:

$$\mathcal{L}_{overall} = \mathcal{L}_{direct} + \lambda\mathcal{L}_{\mathrm{CSS}}, \tag{6}$$

where we use $\lambda = 5$ in our experiments.

4 Experiments

4.1 Experimental Setting

Methodology. We demonstrate the versatility of UniHDA on multi-modal hybrid domain adaptation, *i.e.*, image-image, text-text, and image-text. To show the generator-agnostic nature of UniHDA, we apply it to three well-known generators, *i.e.*, StyleGAN2 [13], Diffusion model [16], and EG3D [4]. Following previous generative domain adaptation literatures [8, 22, 25, 26, 28, 29, 47, 51, 53], we use StyleGAN2 for comparisons in most experiments.

Datasets. We conduct the experiments for a wide range of source and target domains to validate the effectiveness of UniHDA. Following previous work, we consider FFHQ [12], AFHQ-Dog [5], and LSUN-Church [50] as the source domains. The resolutions of images in these datasets are respectively 1024, 512, and 256. We adapt the pre-trained models to diverse hybrid domains driven by the text prompt and one-shot image. To demonstrate the effect of the hybrid domain, *we set the domain coefficients in Eq. (4) as 0.5.*
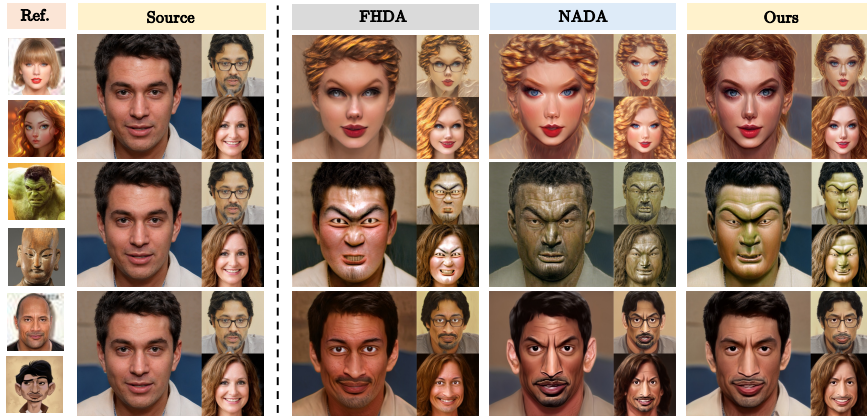
Fig. 5: Image-image hybrid domain adaptation. We compare the results of FHDA [20], NADA [8] and UniHDA (Ours) with the same noise. FHDA and NADA generate images with poor cross-domain consistency, leading to a limited diversity. In contrast, UniHDA alleviates overfitting and maintains strong cross-domain consistency.

Method	Taylor-Elena		Hulk-Wooden		Johnson-Comic		Average	
	CS-I (↑)	SCS (↑)	CS-I (↑)	SCS (↑)	CS-I (↑)	SCS (↑)	CS-I (↑)	SCS (↑)
FHDA	0.685	0.576	0.635	0.659	0.640	0.679	0.630	0.661
NADA	0.684	0.579	0.624	0.575	0.647	0.642	0.628	0.639
Ours	**0.699**	**0.738**	**0.649**	**0.707**	**0.656**	**0.764**	**0.642**	**0.769**

Table 1: Quantitative results for **image-image** domain adaptation. We present the quantitative results corresponding to each case in Fig. 5. To further demonstrate the robustness of our method, we average the results for 25 cases (shown in Appendix).

Evaluation Metrics. One important aspect to evaluate generative domain adaptation is the preservation of domain-specific characteristics. Following [36], we use CLIP Score (CS-T and CS-I) for text-text and image-image adaptation respectively. Concretely, CS-T is measured by the average cosine similarity between prompt and generated images' embedding. CS-I is the average pairwise cosine similarity between CLIP embeddings of generated and real images. Here we use average CS-T or CS-I of multiple domains. For image-text adaptation, we use the average of CS-T and CS-I as the metric (CS).

Another important evaluation is the cross-domain consistency of the source domain. To measure it, we adopt Structural Consistency Score (SCS) [47] to evaluate the spatial structural consistency between source and target generator.

4.2   Image-image Hybrid Domain Adaptation

Fig. 5 shows the qualitative results of baselines and UniHDA for image-image adaptation, starting from the same source domain FFHQ [12] to the combinations of individual domains. As shown in the figure, FHDA [20] suffers from severe model
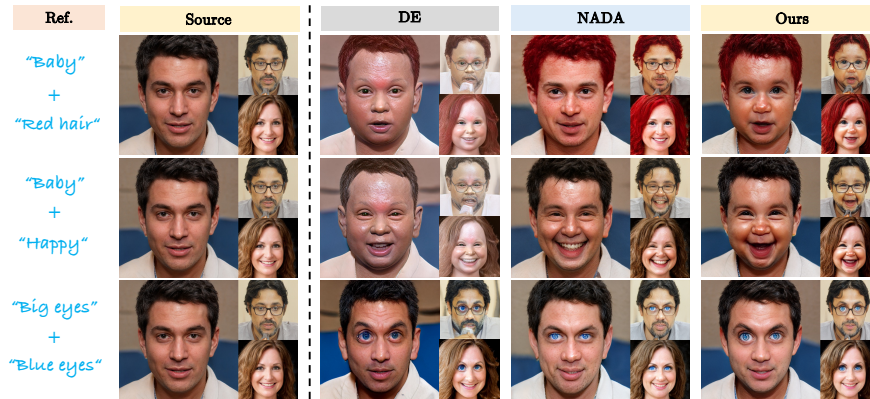
Fig. 6: Text-text hybrid domain adaptation. We compare the results of DE [28], NADA [8] and UniHDA (Ours) with the same noise. UniHDA exhibits desirable performance to acquire characteristics from hybrid target domain and maintain robust cross-domain consistency.

Method	Baby-Red hair		Baby-Happy		Big-Blue eyes		Average	
	CS-T (↑)	SCS (↑)	CS-T (↑)	SCS (↑)	CS-T (↑)	SCS (↑)	CS-T (↑)	SCS (↑)
DE	0.163	0.638	0.160	0.580	0.195	0.662	0.167	0.634
NADA	0.179	0.661	0.170	0.642	0.186	0.731	0.159	0.552
Ours	**0.186**	**0.744**	**0.175**	**0.757**	**0.197**	**0.765**	**0.176**	**0.707**

Table 2: Quantitative results for **text-text** domain adaptation. We present the quantitative results corresponding to each case in Fig. 6. Similar to Tab. 1, we average the quantitative results for 25 cases (shown in Appendix).

collapse and generates images with a limited diversity due to the scarcity of image references. While NADA [8] mitigates overfitting to a certain extent, its cross-domain consistency remains poor, resulting in the generation of similar images. In contrast, UniHDA maintains strong consistency and effectively generates images with characteristics of the hybrid domain.

We also quantitatively compare UniHDA with baselines. As shown in Tab. 1, ours clearly outperforms them. For CS-I, UniHDA significantly outperforms other methods, indicating that generated images effectively integrate multiple characteristics from distinct domains. Furthermore, UniHDA achieves better SCS, which effectively maintains cross-domain consistency compared with baselines.

4.3 Text-text Hybrid Domain Adaptation

Fig. 6 shows the qualitative results for text-text adaptation. Since the adaptation is conducted solely along one projection direction of the latent code, Domain Expansion (DE) [28], to some extent, does not fully capture the characteristics of the target domain, *e.g.*, *baby* (Row 1 and Row 2). Furthermore, DE does not
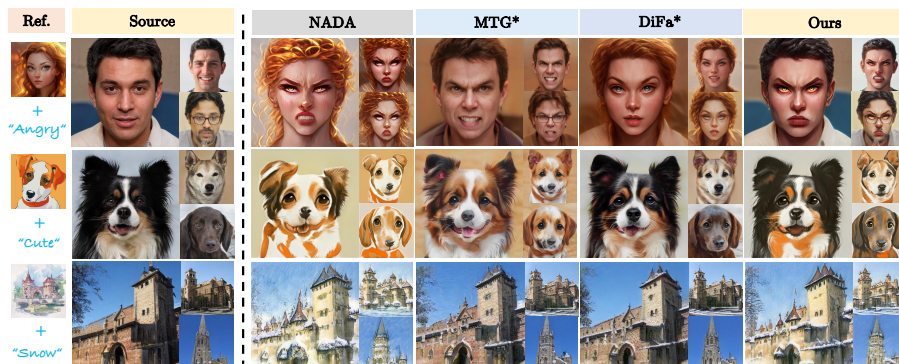
Fig. 7: Image-text hybrid domain adaptation. We compared our method with previous [8, 51, 55], UniHDA well captures the attributes of hybrid target domain and maintains strong cross-domain consistency with source domain. ∗ indicates that MTG and DiFa support multi-modalities by interpolating model parameters with NADA.

Method	FFHQ		Dog		Church	
	CS (↑)	SCS (↑)	CS (↑)	SCS (↑)	CS (↑)	SCS (↑)
NADA	0.563	0.586	0.424	0.533	**0.414**	0.629
MTG	0.536	0.529	0.403	0.526	0.403	0.684
DiFa	0.548	0.681	0.413	0.683	0.407	0.711
Ours	**0.565**	**0.742**	**0.430**	**0.796**	**0.414**	**0.781**

Table 3: Quantitative results for **image-text** adaptation. We average the quantitative results for 81 cases for FFHQ, 16 cases for AFHQ-Dog, and 16 cases for LSUN-Church (shown in Appendix).

Method	Fidel.	Diver.	Corr.
vs. NADA (I-I)	85.2	90.6	76.0
vs. NADA (T-T)	81.4	84.2	80.8
vs. NADA (T-I)	84.6	85.8	78.6

Table 4: User study for fidelity, diversity, and reference corespondence (image or text) in hybrid domain adaptation. The value (%) represents the percentage of users who favor the images generated by our method over NADA.

maintain robust consistency, *e.g.*, the chin of the person in the upper-right corner and background artifacts in Row 3. This is primarily due to the characteristics bias introduced when projecting the latent code to the base subspace, making the images less authentic. The problem of NADA [8] is overfitting. Hard-to-learn characteristics, *e.g.*, *baby* (Row 1 and Row 2) and *big eyes* (Row 3) may be overshadowed by other overfitted ones. In contrast, UniHDA (Ours) exhibits desirable performance to generate images with integrated characteristics while maintaining robust consistency with the source domain.

Similar to Sec. 4.2, we also compare UniHDA with the baselines quantitatively. As shown in Tab. 2, ours clearly outperforms the baselines, which are consistent with qualitative results in Fig. 6. We achieve better CS-I and SCS, indicating that generated images effectively integrate domain-specific attributes and preserve primary characteristics of the source domain.

4.4 Image-text Hybrid Domain Adaptation

Fig. 7 shows the results of image-text adaptation, including FFHQ [12], AFHQ-Dog [5], and LSUN -Church [50]. As depicted in Sec. 4.2, NADA is susceptible to

Method	Modality	Generator Dependency	Model Amount	2-domain		10-domain	
				size ($\downarrow$)	time ($\downarrow$)	size ($\downarrow$)	time ($\downarrow$)
NADA	Multi	-	N	48M	4min	240M	20min
MTG*	Multi	StyleGAN	N	48M	4min	240M	20min
DiFa*	Multi	StyleGAN	N	48M	4min	240M	20min
DE†	Text	Latent space	1	24M	20h	24M	20h
FHDA	Image	-	1	24M	3min	24M	3min
Ours	Multi	-	1	**24M**	**2min**	**24M**	**2min**

Table 5: Comparison with previous methods. ∗ indicates that MTG and DiFa support multi-modalities by interpolating model parameters with NADA. † means that DE needs source dataset (e.g., FFHQ) that significantly increases the training time. Besides, MTG, DiFa, and DE have additional dependencies on the type of generator, which limits their broader applicability.

overfitting, which retains poor cross-domain consistency. Besides, we interpolate NADA's parameters with MTG [55] and DiFa [51], which alleviates overfitting to some extent. However, they can't accurately capture the attributes of hybrid target domain and still fail to maintain good consistency. In contrast, UniHDA well captures the attributes and achieves robust consistency in all scenarios.

As shown in Tab. 3, we also compare UniHDA with the baselines quantitatively. Consistent with qualitative results in Fig. 7, ours clearly outperforms the baselines. Besides, we conduct user study in Tab. 4. For each case, we generate 1000 samples and randomly assign 200 samples to 30 users. The results indicate that UniHDA surpasses NADA in terms of fidelity, diversity and reference correspondence.

4.5   Comparison with Existing Methods

In addition to surpassing baselines for generation quality, UniHDA also surpasses them in terms of efficiency, *e.g.*, model size and training time as shown in Tab. 5. NADA, MTG, and DiFa trains a separate generative model per domain and interpolates their parameters in test-time, which necessitates multiple times the model size and training time. Although DE avoids cross-model interpolation, it heavily relies on the large source dataset (*e.g.*, FFHQ [12]) for regularization during training process, resulting in a significant increase in training time. In contrast, UniHDA circumvents these issues, which enables the completion of the adaptation within single generator in only two minutes.

Furthermore, DE relies on the semantic latent space of the generator (*e.g.*, StyleGAN [12] and DiffAE [33]) for hybrid domain adaptation, limiting its applicability to a broader range of generators. MTG and DiFa utilize GAN inversion, which restricts the applicability to generators similar to StyleGAN. Conversely, UniHDA is not constrained by the type of generators, allowing for its broader application across various generators.

Fig. 8: Hybrid domain adaptation in 3D generator. To show the versatility of UniHDA, we apply it on the popular 3D-aware generator, EG3D [4].

4.6   Generalization on 3D Generator

To further verify the versatility of UniHDA, we apply it on the popular 3D-aware image generation method, EG3D [4]. Specifically, we replace the discrimination loss with our proposed $\mathcal{L}_{direct}$ and $\mathcal{L}_{\mathrm{CSS}}$ for hybrid domain adaptation. As shown in Fig. 8, the results effectively integrate the attributes and preserve the characters and poses of source domain, demonstrating that UniHDA is agnostic to the type of generative models and can be easily extended to other generators.

4.7   Generalization on Diffusion Model

In this section, we demonstrate that UniHDA is agnostic to the type of generative models and can easily generalize to diffusion models. Specifically, we apply UniHDA on DiffusionCLIP [16] and replace the training objective of DiffusionCLIP with our proposed $\mathcal{L}_{direct}$ and $\mathcal{L}_{\mathrm{CSS}}$. As shown in Fig. 9, the results integrate the characteristics from multiple target domains and maintain robust consistency with the source domain. This verifies desirable generalization ability and versatility of UniHDA. More results are included in Appendix.

Fig. 9: Results of UniHDA with DiffusionCLIP [16], which demonstrate UniHDA is agnostic to the type of generator, allowing for broader application on diffusion models.

Fig. 10: Ablation of our proposed $\mathcal{L}_{\text{CSS}}$ on hybrid domain adaptation, which significantly alleviates overfitting and improves cross-domain consistency.

4.8   Ablation of CSS Loss

We conduct the ablation study to evaluate the effects of our proposed Cross-domain Spatial Structure (CSS) loss on single domains. As shown in Fig. 10, the results without $\mathcal{L}_{\text{CSS}}$ suffer from overfitting and have very limited cross-domain consistency, *e.g.*, distorted backgrounds in Row 1 and 3. Benefited from CSS, the generated images maintain consistency with the source images in terms of spatial structure, thereby inheriting the diversity from the source domain. More results are included in Appendix. Besides, we conduct the quantitative ablation in Tab. 6, which is consistent with the qualitative results.

4.9   Ablation of Encoder for CSS

We conduct experiments on pre-trained ViT [6], MViTv2 [21], and Dinov2 to explore the impact of different image encoders for CSS. As shown in Fig. 11, we can observe that all of them improve the consistency with source domain compared with the baseline approach. Furthermore, they exhibit a similar qualitative style, which demonstrates that our CSS is agnostic to different pre-trained image encoders. More results are included in Appendix.
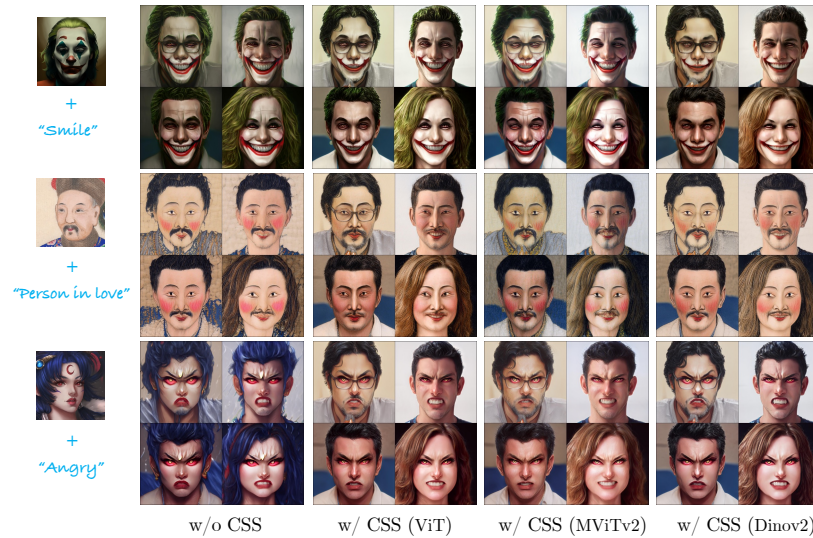
Fig. 11: Ablation of different pre-trained encoders for CSS on hybrid domain adaptation.

λ	FFHQ (I-I)		FFHQ (T-T)		FFHQ (T-I)		Dog (T-I)		Church (T-I)	
	SCS (↑)	CS-I (↑)	SCS (↑)	CS-T (↑)	SCS (↑)	CS (↑)	SCS (↑)	CS (↑)	SCS (↑)	CS (↑)
0	0.502	0.639	0.520	0.170	0.562	0.557	0.491	**0.430**	0.604	0.411
3	0.681	0.638	0.683	0.171	0.694	0.556	0.787	0.428	0.706	0.413
5	**0.769**	**0.642**	**0.707**	**0.176**	**0.742**	**0.565**	**0.796**	**0.430**	**0.781**	**0.414**

Table 6: Quantitative ablation for our proposed $\mathcal{L}_{\text{CSS}}$.

5   Conclusion & Limitation

In this paper, we propose UniHDA, a unified and versatile framework for multi-modal hybrid domain adaptation. To enable multiple modalities, we leverage CLIP encoder to project the references into a unified embedding space. For hybrid domain, we demonstrate the compositional capabilities of direction vectors in CLIP's embedding space and linearly interpolate direction vectors of multiple target domains. In addition, we propose a new cross-domain spatial structure loss to improve consistency, which is conducted in generator-agnostic space and versatile for various generators. We believe our work is an important step towards generative domain adaptation, since we have demonstrated the source generator can be effectively adapted to a hybrid domain with multi-modal references and maintain robust cross-domain consistency. Our code will be made public.

While UniHDA effectively realizes multi-modal hybrid domain adaptation, it also has the limitation. To encode both image and text into a shared embedding space, we utilize pre-trained CLIP during training time, which might bring potential bias for some domains. Nevertheless, we believe that the exploration of the novel task is significant for future work and solutions could be integrated into UniHDA to eliminate the bias.

References

1. Alanov, A., Titov, V., Nakhodnov, M., Vetrov, D.: Styledomain: Efficient and lightweight parameterizations of stylegan for one-shot and few-shot domain adaptation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2184–2194 (October 2023) 2, 4

2. Alanov, A., Titov, V., Vetrov, D.P.: Hyperdomainnet: Universal domain adaptation for generative adversarial networks. Advances in Neural Information Processing Systems **35**, 29414–29426 (2022) 2, 4

3. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. In: ICLR (2018) 2

4. Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., et al.: Efficient geometry-aware 3d generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16123–16133 (2022) 3, 7, 12

5. Choi, Y., Uh, Y., Yoo, J., Ha, J.W.: StarGAN v2: Diverse image synthesis for multiple domains. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8188–8197 (2020) 7, 10

6. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020) 3, 7, 13

7. Duan, Y., Niu, L., Hong, Y., Zhang, L.: Weditgan: Few-shot image generation via latent space relocation. arXiv preprint arXiv:2305.06671 (2023) 2, 4

8. Gal, R., Patashnik, O., Maron, H., Chechik, G., Cohen-Or, D.: Stylegan-nada: Clip-guided domain adaptation of image generators. arXiv preprint arXiv:2108.00946 (2021) 2, 4, 5, 7, 8, 9, 10

9. Härkönen, E., Hertzmann, A., Lehtinen, J., Paris, S.: Ganspace: Discovering interpretable gan controls. Advances in Neural Information Processing Systems **33**, 9841–9850 (2020) 3, 4, 6

10. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Neural Information Processing Systems,Neural Information Processing Systems (Jan 2020) 3, 5

11. Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., Aila, T.: Alias-free generative adversarial networks. Advances in Neural Information Processing Systems **34**, 852–863 (2021) 3, 5

12. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4401–4410 (2019) 2, 3, 4, 5, 7, 8, 10, 11

13. Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8110–8119 (2020) 3, 5, 7

14. Kim, G., Chun, S.Y.: Datid-3d: Diversity-preserved domain adaptation using text-to-image diffusion for 3d generative model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14203–14213 (2023) 2, 4

15. Kim, G., Jang, J.H., Chun, S.Y.: Podia-3d: Domain adaptation of 3d generative model across large domain gap using pose-preserved text-to-image diffusion. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 22603–22612 (2023) 2, 4

16. Kim, G., Kwon, T., Ye, J.C.: Diffusionclip: Text-guided diffusion models for robust image manipulation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2426–2435 (2022) 3, 5, 7, 12, 13

17. Kim, K., Kim, Y., Cho, S., Seo, J., Nam, J., Lee, K., Kim, S., Lee, K.: Diffface: Diffusion-based face swapping with facial guidance. arXiv preprint arXiv:2212.13344 (2022) 2, 4

18. Kwon, M., Jeong, J., Uh, Y.: Diffusion models already have a semantic latent space. arXiv preprint arXiv:2210.10960 (2022) 5

19. Lei, B., Yu, K., Feng, M., Cui, M., Xie, X.: Diffusiongan3d: Boosting text-guided 3d generation and domain adaption by combining 3d gans and diffusion priors. arXiv preprint arXiv:2312.16837 (2023) 2, 4

20. Li, H., Liu, Y., Xia, L., Lin, Y., Zheng, T., Yang, Z., Wang, W., Zhong, X., Ren, X., He, X.: Few-shot hybrid domain adaptation of image generators. arXiv preprint arXiv:2310.19378 (2023) 4, 8

21. Li, Y., Wu, C.Y., Fan, H., Mangalam, K., Xiong, B., Malik, J., Feichtenhofer, C.: Mvitv2: Improved multiscale vision transformers for classification and detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4804–4814 (2022) 13

22. Li, Y., Zhang, R., Lu, J.C., Shechtman, E.: Few-shot image generation with elastic weight consolidation. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) Advances in Neural Information Processing Systems. vol. 33, pp. 15885–15896. Curran Associates, Inc. (2020), https://proceedings.neurips.cc/paper/2020/file/b6d767d2f8ed5d21a44b0e5886680cb9-Paper.pdf 2, 4, 5, 7

23. Liu, Z., Li, L., Xiao, J., Zha, Z.J., Huang, Q.: Text-driven generative domain adaptation with spectral consistency regularization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 7019–7029 (October 2023) 2, 4

24. Lyu, Y., Lin, T., Li, F., He, D., Dong, J., Tan, T.: Deltaedit: Exploring text-free training for text-driven image manipulation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6894–6903 (2023) 2, 4

25. Mo, S., Cho, M., Shin, J.: Freeze the discriminator: a simple baseline for fine-tuning gans. In: CVPR AI for Content Creation Workshop (2020) 2, 4, 5, 7

26. Mondal, A.K., Tiwary, P., Singla, P., Prathosh, A.: Few-shot cross-domain image generation via inference-time latent-code learning. In: The Eleventh International Conference on Learning Representations 2, 4, 5, 7

27. Nitzan, Y., Gharbi, M., Zhang, R., Park, T., Zhu, J.Y., Cohen-Or, D., Shechtman, E.: Domain expansion of image generators (2023) 2, 4

28. Nitzan, Y., Gharbi, M., Zhang, R., Park, T., Zhu, J.Y., Cohen-Or, D., Shechtman, E.: Domain expansion of image generators. arXiv preprint arXiv:2301.05225 (2023) 4, 7, 9

29. Ojha, U., Li, Y., Lu, J., Efros, A.A., Lee, Y.J., Shechtman, E., Zhang, R.: Few-shot image generation via cross-domain correspondence. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10743–10752 (2021) 2, 4, 5, 7

30. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018) 3, 7

31. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. arXiv preprint arXiv:2304.07193 (2023) 3, 7

32. Patashnik, O., Wu, Z., Shechtman, E., Cohen-Or, D., Lischinski, D.: Styleclip: Text-driven manipulation of stylegan imagery. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2085–2094 (2021) 4

33. Preechakul, K., Chatthee, N., Wizadwongsa, S., Suwajanakorn, S.: Diffusion autoencoders: Toward a meaningful and decodable representation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10619–10629 (2022) 4, 5, 11

34. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. pp. 8748–8763. PMLR (2021) 2, 4, 5

35. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10684–10695 (2022) 2

36. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22500–22510 (2023) 7, 8

37. Shen, Y., Yang, C., Tang, X., Zhou, B.: Interfacegan: Interpreting the disentangled face representation learned by gans. IEEE transactions on pattern analysis and machine intelligence 44(4), 2004–2018 (2020) 4

38. Shen, Y., Zhou, B.: Closed-form factorization of latent semantics in gans. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1532–1540 (2021) 3, 4, 6

39. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv: Learning,arXiv: Learning (Oct 2020) 5

40. Song, K., Han, L., Liu, B., Metaxas, D., Elgammal, A.: Diffusion guided domain adaptation of image generators. arXiv preprint arXiv:2212.04473 (2022) 2, 4

41. Spingarn-Eliezer, N., Banner, R., Michaeli, T.: Gan" steerability" without optimization. arXiv preprint arXiv:2012.05328 (2020) 4

42. Tov, O., Alaluf, Y., Nitzan, Y., Patashnik, O., Cohen-Or, D.: Designing an encoder for stylegan image manipulation. ACM Transactions on Graphics (TOG) 40(4), 1–14 (2021) 4

43. Vahdat, A., Kreis, K., Kautz, J.: Score-based generative modeling in latent space. Advances in Neural Information Processing Systems 34, 11287–11302 (2021) 2

44. Voynov, A., Babenko, A.: Unsupervised discovery of interpretable directions in the GAN latent space. arXiv preprint arXiv:2002.03754 (2020) 4

45. Wu, Y., Li, Z., Wang, C., Zheng, H., Zhao, S., Li, B., Ta, D.: Domain re-modulation for few-shot generative domain adaptation. arXiv preprint arXiv:2302.02550 (2023) 2, 4

46. Wu, Z., Lischinski, D., Shechtman, E.: StyleSpace analysis: Disentangled controls for StyleGAN image generation. arXiv:2011.12799 (2020) 4

47. Xiao, J., Li, L., Wang, C., Zha, Z.J., Huang, Q.: Few shot generative model adaption via relaxed spatial structural alignment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11204–11213 (2022) 2, 4, 5, 7, 8

48. Xu, G., Hou, Y., Liu, Z., Loy, C.C.: Mind the gap in distilling stylegans. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII. pp. 423–439. Springer (2022) 3, 4, 6

49. Ye, H., Zhang, J., Liu, S., Han, X., Yang, W.: Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. arXiv preprint arXiv:2308.06721 (2023) 2, 4
50. Yu, F., Seff, A., Zhang, Y., Song, S., Funkhouser, T., Xiao, J.: Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. arXiv preprint arXiv:1506.03365 (2015) 7, 10
51. Zhang, Y., mingshuai Yao, Wei, Y., Ji, Z., Bai, J., Zuo, W.: Towards diverse and faithful one-shot adaption of generative adversarial networks. In: Oh, A.H., Agarwal, A., Belgrave, D., Cho, K. (eds.) Advances in Neural Information Processing Systems (2022), https://openreview.net/forum?id=IXoHxXIGpyV 2, 4, 5, 7, 10, 11
52. Zhao, Y., Chandrasegaran, K., Abdollahzadeh, M., Cheung, N.M.M.: Few-shot image generation via adaptation-aware kernel modulation. Advances in Neural Information Processing Systems 35, 19427–19440 (2022) 2, 4
53. Zhao, Y., Ding, H., Huang, H., Cheung, N.M.: A closer look at few-shot image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9140–9150 (2022) 2, 4, 5, 7
54. Zhao, Y., Du, C., Abdollahzadeh, M., Pang, T., Lin, M., Yan, S., Cheung, N.M.: Exploring incompatible knowledge transfer in few-shot image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7380–7391 (2023) 2, 4
55. Zhu, P., Abdal, R., Femiani, J., Wonka, P.: Mind the gap: Domain gap control for single shot domain adaptation for generative adversarial networks. arXiv preprint arXiv:2110.08398 (2021) 2, 4, 10, 11
56. Zhu, Y., Liu, H., Song, Y., Yuan, Z., Han, X., Yuan, C., Chen, Q., Wang, J.: One model to edit them all: Free-form text-driven image manipulation with semantic modulations. Advances in Neural Information Processing Systems 35, 25146–25159 (2022) 2, 4

UniHDA: A Unified and Versatile Framework for Multi-Modal Hybrid Domain Adaptation Supplementary Material

Hengjia Li[1], Yang Liu[1], Yuqi Lin[1], Zhanwei Zhang[1], Yibo Zhao[1], Weihang Pan[1], Tu Zheng[2], Zheng Yang[2], Chunjiang Yu[3], Boxi Wu[1], and Deng Cai[1]

[1] Zhejiang University   [2] Fabu Inc.   [3] Ningbo Port

Fig. 1: Hybrid domain adaptation from AFHQ *cat* to incompatible domains, *i.e. lion*, *rabbit*, and *tiger*.

1 Appendix

1.1 Results of Incompatible Domain Adaptation

To verify the generalization ability of UniHDA, we additionally conduct experiments for hybrid domain adaptation on incompatible domains, *i.e.* from *cat* to *rabbit*. As shown in Fig. 1 and Fig. 2, we start from AFHQ [1] *cat* and *dog* to incompatible domains respectively. We can observe that the results effectively integrate the attributes of the corresponding domain and maintain robust consistency with source domain.

1.2 More Qualitative Results

We apply UniHDA to adapt the generator on FFHQ [3] to more hybrid domains, *i.e.*, text-text, image-image, and image-text. As shown in Fig. 3, Fig. 4, Fig. 5,
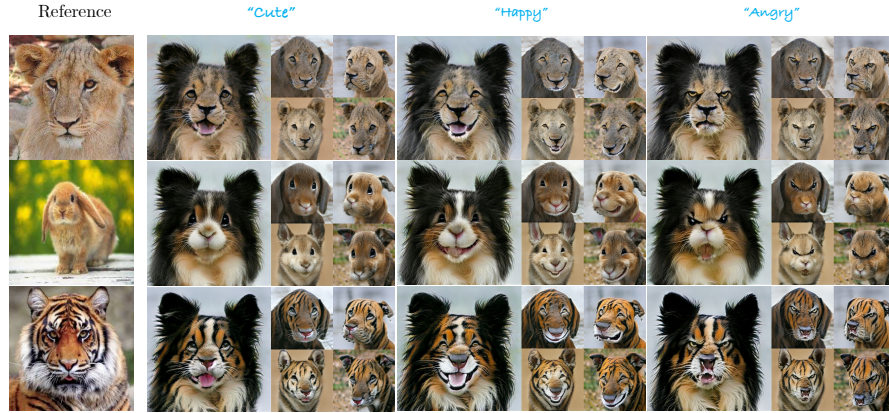
Fig. 2: Hybrid domain adaptation from AFHQ *dog* to incompatible domains, *i.e. lion*, *rabbit*, and *tiger*.

Fig. 3: The results of hybrid domain adaptation from FFHQ to the hybrid of more domains.

and Fig. 6, UniHDA successfully generates images with integrated characteristics from multiple target domains and maintains robust consistency with the source domain. Besides, we showcase more results of hybrid domain adaptation from AFHQ *dog* and LSUN *church* [5] in Fig. 7 and Fig. 8.

1.3   More Ablation of CSS Loss

As depicted in Sec. 4.8 of the main paper, our proposed $\mathcal{L}_{\text{CSS}}$ significantly alleviates overfitting and improves cross-domain consistency. Results with $\mathcal{L}_{\text{CSS}}$ achieve better SCS score, indicating that they maintain stronger consistency with the source domain. Additionally, we show more qualitative results in Fig. 9 to verify the effectiveness of UniHDA.

1.4   More Ablation of Encoder for CSS

In Sec. 4.9 of the main paper, we conduct experiments on pre-trained ViT, MViTv2, and Dinov2 to explore the impact of different image encoders for CSS.

Fig. 4: More results of image-image hybrid domain adaptation. The source image is in the top-left corner, and the first row and column consist of training images.

Additionally, we showcase the results of image-image and text-text in Fig. 10 and Fig. 11 respectively. We can observe that all of them improve the consistency with source domain compared with the baseline approach. Furthermore, they exhibit a similar qualitative style, which demonstrates that our CSS is agnostic to different pre-trained image encoders.

1.5   More Results for DiffusionCLIP

To demonstrate the versatility of UniHDA, we apply it on DiffusionCLIP in Sec. 4.7 in the main paper. As shown in Fig. 12, we showcase more results for DiffusionCLIP including image-image, text-text, and image-text. All results achieve hybrid domain adaptation and preserve strong corss-domain consistency.
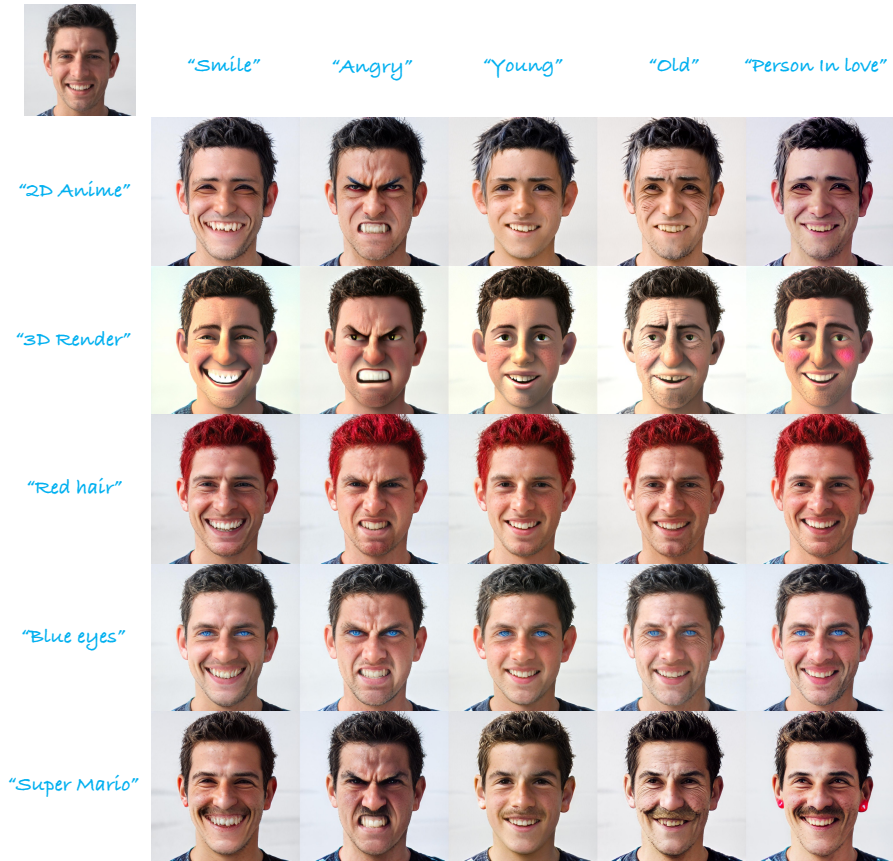
Fig. 5: More results of text-text hybrid domain adaptation. The source image is in the top-left corner, and the first row and column consist of text prompts.

1.6   Implement Details

Following the setting of previous generative domain adaptation methods [2, 4], we utilize the batch size of 4 and ADAM Optimizer with a learning rate of 0.002 for all experiments during training. A training session typically requires 300 iterations in 2 minutes, which significantly reduces training time compared with adversarial methods for generative domain adaptation. Note that we conduct all experiments on a single NVIDIA RTX 4090 GPU. The code will be open source.

For experiments on FFHQ, we generate images with 1024×1024 resolution. As for AFHQ-Dog and LSUN-Church, we operate on 512×512 and 256×256 resolution images respectively.

References

1. Choi, Y., Uh, Y., Yoo, J., Ha, J.W.: StarGAN v2: Diverse image synthesis for multiple domains. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8188–8197 (2020) 1
2. Gal, R., Patashnik, O., Maron, H., Chechik, G., Cohen-Or, D.: Stylegan-nada: Clip-guided domain adaptation of image generators. arXiv preprint arXiv:2108.00946 (2021) 4
3. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4401–4410 (2019) 1
4. Nitzan, Y., Gharbi, M., Zhang, R., Park, T., Zhu, J.Y., Cohen-Or, D., Shechtman, E.: Domain expansion of image generators. arXiv preprint arXiv:2301.05225 (2023) 4
5. Yu, F., Seff, A., Zhang, Y., Song, S., Funkhouser, T., Xiao, J.: Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. arXiv preprint arXiv:1506.03365 (2015) 2
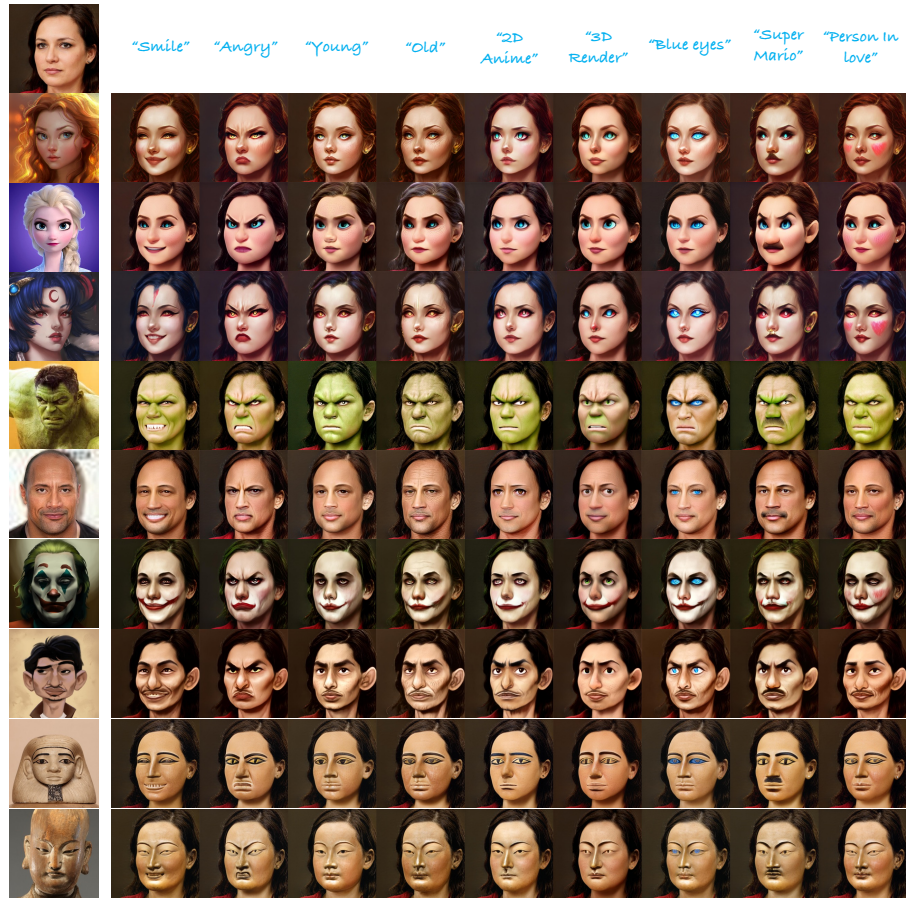
Fig. 6: More results of image-text hybrid domain adaptation. The source image is in the top-left corner. The first row and column consist of training images and text prompts respectively.

Fig. 7: More results of image-text hybrid domain adaptation on AFHQ *dog*.

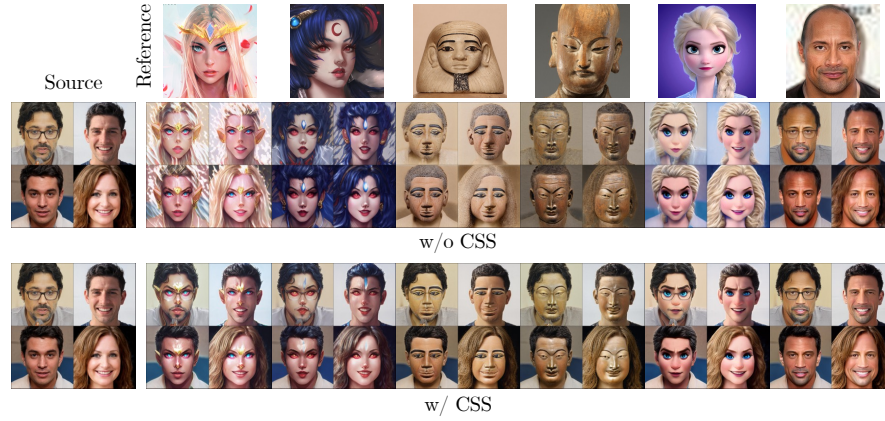**Fig. 8:** More results of image-text hybrid domain adaptation on LSUN *church*.

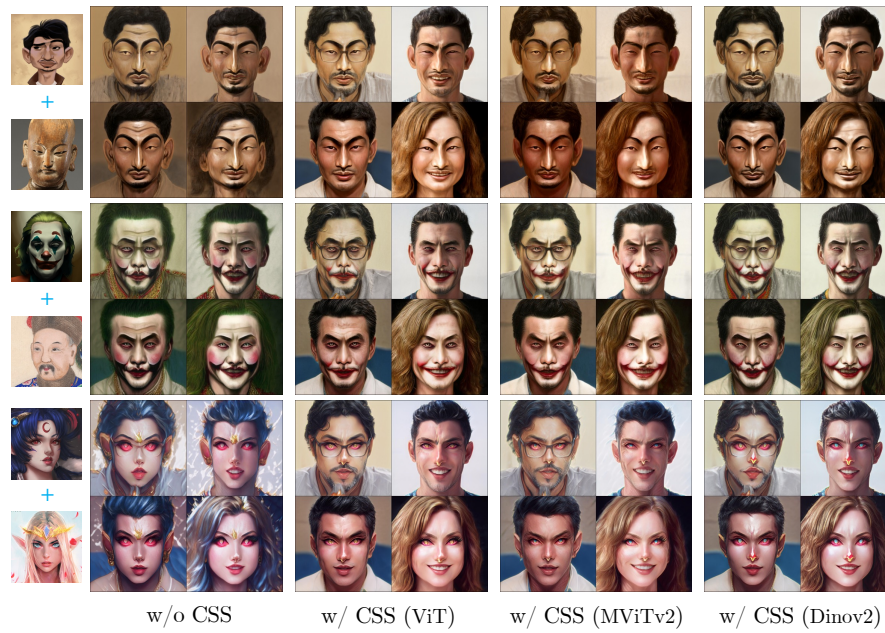**Fig. 9:** More qualitative results to verify the effectiveness of our proposed $\mathcal{L}_{\text{CSS}}$.

w/o CSS   w/ CSS (ViT)   w/ CSS (MViTv2)   w/ CSS (Dinov2)

Fig. 10: Effect of different pre-trained image encoders for CSS on image-image hybrid domain adaptation.

"Angry" + "3D Render"

"Young" + "Blue eyes"

"Smile" + "Red hair"

w/o CSS   w/ CSS (ViT)   w/ CSS (MViTv2)   w/ CSS (Dinov2)
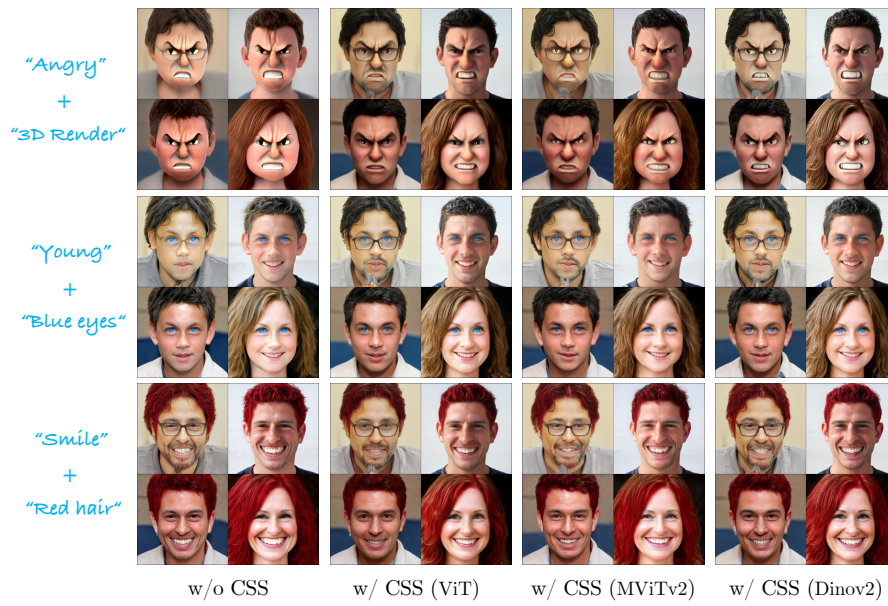
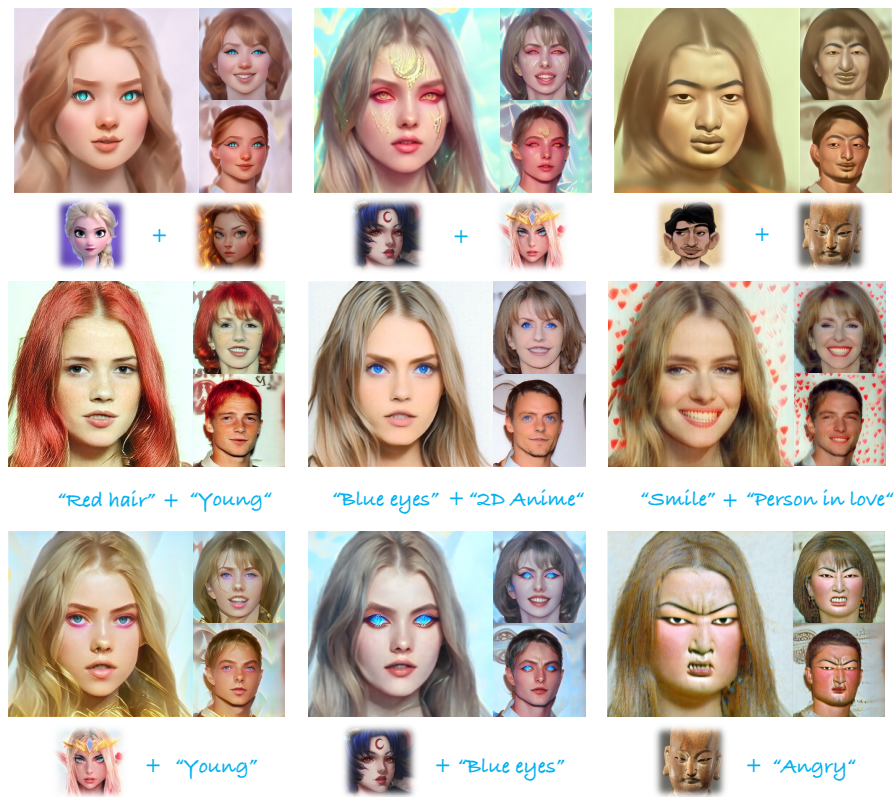**Fig. 11:** Effect of different pre-trained image encoders for CSS on text-text hybrid domain adaptation.

Fig. 12: More results of UniHDA with DiffusionCLIP.