

Separating Content and Style for Unsupervised Image-to-Image Translation

Yunfei Liu¹

lyunfei@buaa.edu.cn

Haofei Wang²

wanghf@pcl.ac.cn

Yang Yue³

yuey23@buaa.edu.cn

Feng Lu^{1,2, ✉}

lufeng@buaa.edu.cn

¹ State Key Laboratory of Virtual Reality Technology and Systems, School of Computer Science and Engineering, Beihang University, Beijing, China

² Peng Cheng Laboratory, Shenzhen, China

³ Beihang University, Beijing, China

Abstract

Unsupervised image-to-image translation aims to learn the mapping between two visual domains with unpaired samples. Existing works focus on disentangling domain-invariant content code and domain-specific style code individually for multimodal purposes. However, less attention has been paid to interpreting and manipulating the translated image. In this paper, we propose to separate the content code and style code simultaneously in a unified framework. Based on the correlation between the latent features and the high-level domain-invariant tasks, the proposed framework demonstrates superior performance in multimodal translation, interpretability and manipulation of the translated image. Experimental results show that the proposed approach outperforms the existing unsupervised image translation methods in terms of visual quality and diversity. Code and data have been released at <https://github.com/DreamtaleCore/SCS-UIT>.

1 Introduction

Unsupervised image-to-image translation (UIT) attracts great attention since it can learn the mapping between different visual domains without paired data. Numerous computer vision and graphics problems can be formulated as UIT problems, such as image super-resolution [6, 24, 50], in-painting [68], style-transfer [18], and other low-level vision tasks [29].

Learning the mappings between two visual domains are inherently multimodal, *i.e.*, a single input may correspond to multiple possible outputs. For the source visual domain \mathcal{X}_A and target domain \mathcal{X}_B , given the source image $x_A \in \mathcal{X}_A$, UIT aims to keep the content of x_A , and to turn its style into the target domain. However, It is not easy to define the content and style precisely. As with the existing methods, we make similar assumptions [16, 25] for the content code and style code. In general, the content is shared between two domains and

often refers more to semantics, layouts, and spatial arrangements of the image, while the style involves more in color, tones, and textures, and other domain-specific features.

To achieve such translation, several recent works using cycle consistency [28, 61] and disentangled representations have been proposed. For example, MUNIT [16] and DRIT [25] encode an input image x into two different features, the style feature s and the content feature c , as illustrated in Fig. 1(a). In their settings, domain \mathcal{X}_A and \mathcal{X}_B have their own domain-specific feature spaces for styles \mathcal{S}_A and \mathcal{S}_B . Meanwhile, they share the same latent domain-invariant space \mathcal{C} for the content. Thus, given a source image x_A and a target image x_B , they are encoded as (c_A, s_A) and (c_B, s_B) , respectively. The content codes c_A and c_B belong to \mathcal{C} , while style codes s_A and s_B belong to \mathcal{S}_A and \mathcal{S}_B , respectively. For cross domain translation task $\mathcal{X}_A \mapsto \mathcal{X}_B$, the content code c_A and the style code s_B are combined together for synthesizing the result with content from x_A and style from x_B .




Figure 1: Comparison of the existing approaches on unsupervised image-to-image translation and our proposed method. (a) MUNIT [16], DRIT [25] and DSMAP [4] extract the shared content code and the domain-specific style code through two independent encoders. (b) The proposed SCS-UIT framework separates the content code and style code in feature space on the basis of correlation to the domain-invariant high-level vision task.

Recently, several approaches have been proposed to learn the content code with additional high-level tasks under supervision. INIT [43] uses an additional model to predict the bounding boxes of the instances in the image, and applies both instance-level style and global style to the target image spatially. TransGaGa [52] uses landmarks as the shared geometry information and translates images from two visual domains with apparent geometry gap. However, these works mainly focus on the translation between two domains, while the interpretation and manipulation during the translation have not been well-explored.

In this work, we proposed the Separating Content and Style for Unsupervised Image-to-image Translation (SCS-UIT) framework, which extracts and separates content and style in a unified feature space. More specifically, we first define the Domain-Invariant High-level Vision task (DI-HV) for UIT task. Then we separate style code and content code from different feature channels on the basis of the correlation to DI-HV. The proposed framework has several good properties: 1) it supports multimodal translation within a single encoder, 2) it achieves good interpretability and 3) the translated results can be easily manipulated. In summary, the primary contributions of this work are as follows:

- We introduced the Domain-Invariant High-level Vision (DI-HV) task for unsupervised image-to-image translation. Based on the correlation to DI-HV, a new unsupervised image-to-image translation framework (SCS-UIT) is proposed.
- The SCS-UIT not only eases the interpretation of the translation but also enables us to

interactively edit the semantic regions in the translated image.

- Experimental results demonstrate that our method outperforms the state-of-the-art methods on four image translation tasks.

2 Related Work

Image-to-image translation aims at learning the mapping from the source domain to the target domain. Pix2pix [17] is the first work using conditional GANs for image-to-image translation. Following [17], several works seek to address other computer vision tasks, such as super-resolution [24, 50], domain adaption [13, 26], colorization [8, 55] and low-light enhancement [31, 52].

However, paired data is not always available in practical applications. CycleGAN [50] and UNIT [28] are trained with unpaired data, which equips cycle consistency for UIT. DiscoGAN [22], DualGAN [64] and SelectionGAN [47] are proposed following the idea of cycle-consistency. Other works apply attention mechanism to UIT [7, 21, 35, 53].

Some works attempt to address one-to-many translation or many-to-many translation. By assuming that an image can be decomposed into a domain-invariant content code and a domain-specific style code, CIIT [27], EG-UNIT [63], MUNIT [16] and DRIT [25] learn a one-to-many mapping (*i.e.* multimodal) between the two image domains in an unsupervised settings. Other works add additional constrains to improve style-content disentanglement in image-to-image translation [4, 8, 37, 40, 42, 59]. More recently, high-level vision tasks such as object landmark detection [52], instance bound box detection [43] and semantic parsing [41, 51] are used for translation. However, they usually require an specifically designed architecture for these tasks. Unlike the aforementioned methods, we propose a correlation-based framework to achieve multimodal UIT, which also enables us to interpret the translation process and manipulate the translation result.

Interpreting the learned CNNs helps us to gain more insights of the network design. There are two branches, one branch is to explain individual network decisions using informative feature maps of instances [39]. Another branch is to explore the activation of units in CNNs using modified back-propagation [45, 48]. Morcos *et al.* [36] examined the effect of individual units by an ablation study. Bau *et al.* [9] visualized the GANs by manipulating specific channels in the latent space of the image. To manipulate the content in image translation, disentangling factors of image variations has attracted much attention [8, 19, 20]. These works inspire us to interpret and manipulate the image-to-image translation results.

3 Methodology

In this paper, we propose to separate content and style for unsupervised image-to-image translation (SCS-UIT). This framework reveals the learned relationship of representation between content and style and separates them within one encoder. Therefore, we can interpret and manipulate the translated images semantically based on the learned model. Before diving into the design of the SCS-UIT, we first give a definition of the domain-invariant high-level vision task, then give an overview of our method followed by the details.




Figure 2: Comparisons between Domain-Invariant High-level Vision (DI-HV) tasks and Non-DI-HV tasks on the Colored-MNIST dataset. (a) DI-HV are the tasks sharing between different visual domains. (b) Non DI-HV pay more attention to the domain-specific features.

3.1 Domain-invariant high-level vision task

We denote \mathcal{T}_A as the set of all the possible vision tasks that can be performed on \mathcal{X}_A , and \mathcal{T}_B as the set of all the tasks for \mathcal{X}_B . Domain-invariant high-level vision task (DI-HV) is defined as $\mathcal{T}_A \cap \mathcal{T}_B$. For an instance, Fig. 2 illustrates a typical example of DI-HV. In this case, the DI-HV is the digital number classification task, as though the color tones are different between two domains, which can be taken as a non DI-HV task. DI-HV can be easily extended to other image-to-image translations, *e.g.*, face parsing could be a DI-HV for man \leftrightarrow woman translation or cat face \leftrightarrow human face, scene semantic segmentation could be a DI-HV for season translation or CG image \leftrightarrow real image, number classification for ColoredMNSIT, *etc.*.

3.2 Method overview

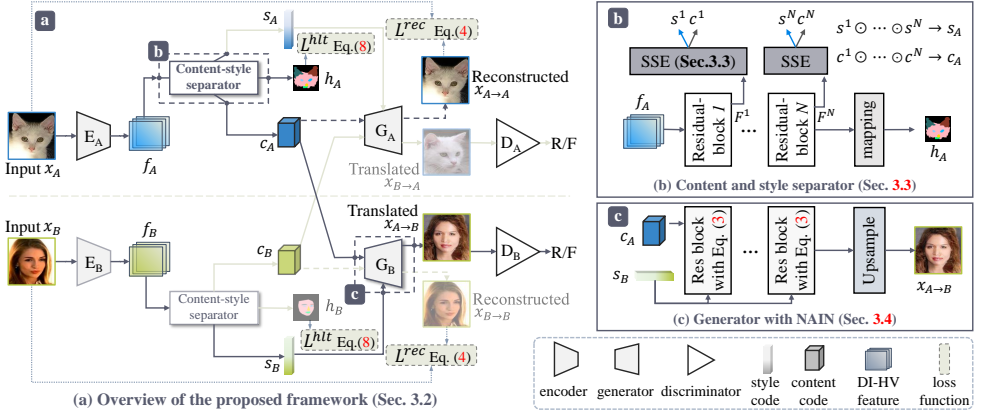


Figure 3: System architecture. (a) Overview of the proposed SCS-UIT framework for UIT. Take cat (x_A) to human ($x_{A \rightarrow B}$) as a example, x_A is first mapped into feature f_A through encoder E_A . Then we separate content and style based on the correlation between the feature and DI-HV task, *i.e.*, facial parsing. Next we synthesize the image using a generator with NAIN, which combines the content code c_A and the style code s_B together to generate the $x_{A \rightarrow B}$. (b) The structure of content-style separator. Note that \odot is channel-wise concatenation. (c) The structure of generator with NAIN.

Fig. 3 (a) demonstrates the overview of our proposed method. In addition to the basic encoders $\{E_A, E_B\}$ for DI-HV, the generators $\{G_A, G_B\}$ and the discriminators $\{D_A, D_B\}$, our method also learns the correlation-based feature separation functions $\Pi_{F_A \rightarrow C, S_A, \mathcal{H}}$ and $\Pi_{F_B \rightarrow C, S_B, \mathcal{H}}$, which 1) separates the DI-HV features into the domain-invariant content space \mathcal{C} and domain-specific style space \mathcal{S} , and 2) maps the feature space into the DI-HV task space \mathcal{H} . For the mapping of $\mathcal{X}_A \mapsto \mathcal{X}_B$, the source image $x_A \in \mathcal{X}_A$ is first encoded into the

DI-HV feature: $f_A = E_A(x_A)$. Next, the content code c_A and style code s_A are separated through $c_A, s_A = \Pi_{F_A \rightarrow C, S_A}(f_A)$. Then the generator G_B takes the content code c_A and the style code s_B , which is sampled from S_B , for translating the output image $x_{A \rightarrow B}$. In summary, for the translation of $\mathcal{X}_A \mapsto \mathcal{X}_B$, our method synthesizes the output by

$$x_{A \rightarrow B} = G_B(\Pi_{F_A \rightarrow C_A}(E_A(x_A)), s_B). \quad (1)$$

Similarly, the translation of $\mathcal{X}_B \mapsto \mathcal{X}_A$ is achieved by

$$x_{B \rightarrow A} = G_A(\Pi_{F_B \rightarrow C_B}(E_B(x_B)), s_A). \quad (2)$$

3.3 Learning to separate content and style

The critical part is to extract content code and style code from the feature space with $\Pi_{F_A \rightarrow C, S_A}$ and $\Pi_{F_B \rightarrow C, S_B}$, and to map features to DI-HV task space. As illustrated in Fig. 3 (b), the content-style separator consists of several residual blocks and a DI-HV task mapping function, where each of them is connected with a Squeeze-Selection-and-Excitation (SSE) module.

Squeeze-Selection-and-Excitation for Content-style Separation. The SSE module is inspired by the SE-Networks [14], which learns weights for different feature channels. The weights describe how important the channels contribute to the final task. In this paper, we take the weight as a descriptor of the correspondence between the feature and DI-HV task. We denote each channel of the feature as a *unit*. Units with high correspondence scores are taken as content code, and the rest of units are taken as style code. In detail, 1) we *squeeze* the units in feature f into channel descriptors via an adaptive average pooling layer. Following [14], we compute the channel weights by learning a non-linear mapping. The weights are taken as the correlation score of each unit. 2) We *select* the units with high correlation scores as content code c , and the rest of units as style code s . 3) The correlation scores then *excite* the units in the feature f for DI-HV task. The input of the SSE module is the feature F_k , which is encoded from the input image. Then this module separates the content code and style code from F_k . The detailed steps are shown in the supplement materials (Algorithm 1).

Flexible mappings for DI-HV tasks. We map the learned features to the DI-HV tasks using supervised learning. The decoder is flexible for different DI-HV tasks. For example, if the image is from domain-invariant dataset (e.g. MNIST), DI-HV task decoder is implemented as a fully connected layer for classification; if the scene semantic segmentation map is available, the high-level task mapping is constructed with a fully convolutional network.

3.4 Image generation with NAIN

In our experiments, we found that the water-drop effects often appeared in the synthesized image. The water-drop effects result in the quality degradation of the generated images. To remove such effects, we propose a normalized adaptive instance normalization (NAIN) block to improve the quality of the generated image.

As reported in StyleGAN2 [20], the effects are mainly caused by the abnormal bias in Adaptive Instance Normalization (AdaIN) [15], which is used in generator for fusing content code and style code. Different from Karras *et al.* [20], where they added a ‘demodulation’ operation, we propose a solution within a unified layer, *i.e.*, Normalized AdaIN (NAIN), to restrain the intermediate parameters in AdaIN. To be more specific, we normalize the

parameters of AdaIN before assigning them to the instance normalization layer. The NAIN is defined as

$$\text{NAIN}(z, \gamma, \beta) = \sigma(\gamma) \left(\frac{z - \mu(z)}{\sigma(z)} \right) + \sigma(\beta), \quad (3)$$

where $\sigma(z) = 1/(1 + \exp(-z))$, z is the activation of the previous convolutional layer, μ and σ are channel-wise mean and standard deviation, γ and β are parameters generated by the multi-layer perceptron, which takes style code as input. NAIN prevents the intermediate variable z from too large or too small values.

3.5 Loss functions

There are four loss functions in our model: image reconstruction loss \mathcal{L}^{rec} , adversarial loss \mathcal{L}^{adv} , perceptual loss $\mathcal{L}^{\text{perc}}$ and DI-HV task loss \mathcal{L}^{hlt} .

Once the input image is reconstructed through auto-encoder with supervised learning, we compute the reconstruction loss as:

$$\mathcal{L}^{\text{rec}} = |x_A - x_{A \rightarrow A}|_1 + |x_B - x_{B \rightarrow B}|_1. \quad (4)$$

To stabilize and accelerate the the GAN training, we use the LSGAN objective proposed by Mao *et al.* [34]. The least-squares adversarial losses [34] are defined as follows:

$$\mathcal{L}_{\mathcal{X}_A}^{\text{adv}} = (1 - D_A(x_A))^2 + (D_A(x_{B \rightarrow A}))^2, \quad \mathcal{L}_{\mathcal{X}_B}^{\text{adv}} = (1 - D_B(x_B))^2 + (D_B(x_{A \rightarrow B}))^2. \quad (5)$$

The total adversarial loss is $\mathcal{L}^{\text{adv}} = \mathcal{L}_{\mathcal{X}_A}^{\text{adv}} + \mathcal{L}_{\mathcal{X}_B}^{\text{adv}}$.

We also add domain-invariant perceptual loss to make the results more realistic. The perceptual loss, often computed as a distance in the VGG [14] feature space between the output and the reference image, accelerates training on high-resolution datasets. The perceptual loss is defined as below:

$$\mathcal{L}^{\text{perc}} = |\Psi(x_A) - \Psi(x_{A \rightarrow B})|_1 + |\Psi(x_B) - \Psi(x_{B \rightarrow A})|_1, \quad (6)$$

where $\Psi(x)$ is the high-level feature of input image x in VGG. We follow the settings of the perceptual loss from MUNIT [17] in our experiments.

For DI-HV task, if the semantic segmentation or classification is applied, we use cross entropy as the loss function:

$$\mathcal{L}^{\text{hlt}} = \sum_{j=1}^K \sum_{i=1}^M (-y_{ji} \log \hat{y}_{ji} - (1 - y_{ji}) \log (1 - \hat{y}_{ji})), \quad (7)$$

where K is the batch size and M is the summation over classes, \hat{y} is the prediction and the ground truth label is y .

Total loss. By using the GAN scheme, we jointly train the encoders E , content-style separator Π , decoders G and discriminators D to optimize the weighted sum of the different loss terms:

$$\min_{E, \Pi, G} \max_D (E, \Pi, G, D) = \mathcal{L}^{\text{adv}} + \lambda_1 * \mathcal{L}^{\text{rec}} + \lambda_2 * \mathcal{L}^{\text{perc}} + \lambda_3 * \mathcal{L}^{\text{hlt}}, \quad (8)$$

where $\lambda_1, \lambda_2, \lambda_3$ are hyper-parameters for adjusting the weights of each loss function.

4 Experiments

Datasets. We conducted extensive experiments on four datasets.

Cat \rightleftharpoons Human is obtained from the Kaggle ¹ and CelebA [60], which are publicly available. For each population of images in different domains, we collected 6000 images (5000 for training and 1000 for testing). Semantic segmentation is taken as the DI-HV. Specifically, we generated the pseudo semantic segmentation mask based on the cat/human facial landmark location.

Man \rightleftharpoons Woman dataset is generated by randomly sampling 10000 images (9000 for training and 1000 for testing) from CelebA [60] for each domain. The semantic segmentation mask is generated in the same way as *Cat \rightleftharpoons Human*.

Colored-MNIST dataset is an extension of MNIST dataset [23]. We followed [9] and generated images with colored digits on black background or white digits on colored background. The image classification is used for the DI-HV.

CG \rightleftharpoons Real dataset is proposed by [9], which translates between synthetic indoor images and real indoor scenes. Semantic scene parsing is used as the DI-HV, where we generate the pseudo semantic segmentation mask by HRNet [49]. The HRNet is trained on ADE20K [60].

We have also tested our methods on other UIT tasks provided by [46]. These tasks include *Cat \rightleftharpoons Dog* (face parsing is used as DI-HV), and *Summer \rightleftharpoons Winter* (semantic segmentation [49] is used as DI-HV).

Evaluation metrics. We evaluated both the realism and diversity of the generated images for different methods. We adopted the Fréchet Inception Distance [42] (FID) to measure visual quality, which is a general image quality assessment method. To measure diversity, we used the Learned Perceptual Image Patch Similarity (LPIPS) to calculate the diversity among images, similar to [53]. Both metrics use InceptionV3 [46] as the backbone, which is pretrained on ImageNet.

Implementation details. Inspired by MUNIT [46], we constructed the SCS-UIT using a similar decoder and discriminator. Images of *Cat \rightleftharpoons Human*, *Man \rightleftharpoons Woman*, and *CG \rightleftharpoons Real* dataset are randomly cropped and resized to 256×256 before feeding into the model. In Algorithm 1, we set select ratio $r = 0.25$. We used 4 residual blocks for the content-style separator. For DI-HV mapping, we applied 2 convolutional blocks with an instance normalization layer for semantic estimation. To optimize the training complexity and performance, we empirically set multi-scale to 3 in the discriminator. We set the iteration to $1e8$ in the experiments. Since images in *Colored-MNIST* are in low resolution, we resized them to 64×64 and use 2 residual blocks for content-style separator. A two-layer MLP is adopted for classification. Other hyper-parameters are set the same as other datasets. We empirically set $\lambda_1 = 5, \lambda_2 = 0.5, \lambda_3 = 1$ for the objective in Eq. (8).

4.1 Interpretation of the learned representations

The proposed framework is able to interpret the translation results. Inspired by the image classification task [42] and recent GAN dissection [9], we find the feature maps, which are from different channels of the our learned content code, represent different semantic information of the input image. Based on this observation, we further manipulate the translations with semantic meanings through operations in feature space.

¹<https://www.kaggle.com/crawford/cat-dataset>

We extract the content codes (64 units) of the input image and find the valid units by calculating the IoUs between the unit and semantic segmentation. Intuitively, if a unit has high overlap with a specific semantic map (*e.g.*, facial region), it will be a valid unit for the facial semantics. Mathematically, the valid unit is defined as follows:

$$U_{\text{valid}}^m = \begin{cases} 1, & \text{if } \text{IoU}(U^*, M) > \alpha; \\ 0, & \text{otherwise.} \end{cases} \quad (9)$$

where $U_{i,j}^* = 1$, if $U_{i,j} > \beta$; otherwise, we set $U_{i,j}$ to 0. α and β are thresholds for IoU and activated units U^* . M is the semantic map, which is used as ground-truth for DI-HV task. We set $\alpha = 0.18$, $\beta = 0.4$ in the experiment. Fig. 4 demonstrates our method is able to extract more semantically valid units in the content code across different image-to-image translation tasks. In each task, we tested 100 images on visualizing the relationship between the semantic masks and the content codes.



Figure 4: Comparison of the number of valid unit in content code between different methods.

4.2 Semantic area manipulation

The proposed framework is capable of interactive manipulation of the semantic areas in the translated images. Here, we conducted two experiments: replacing the semantic area and ablating the artifacts.

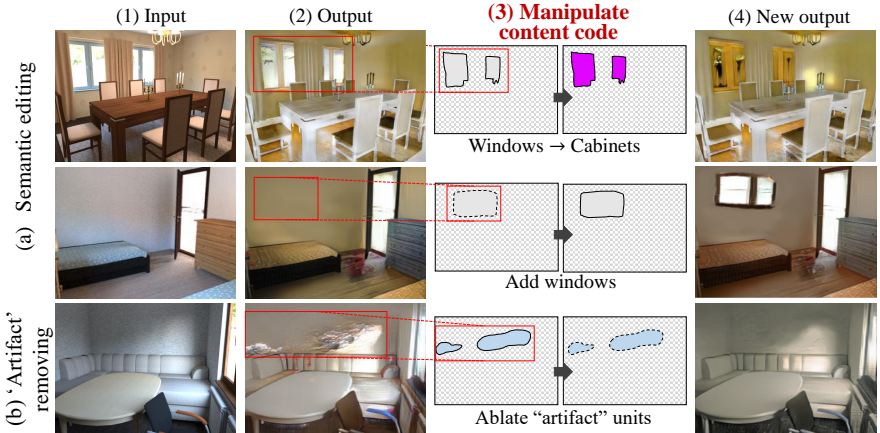


Figure 5: Editing results on the CG \rightarrow Real dataset. The proposed method is capable of (a) editing the semantic regions, such as replacing the windows to cabinets, adding windows on the wall, and (b) removing the ‘artifact’ units.

We first show the results of editing the semantic area. Take the first row of Fig. 5 as an example, 1) we locate the IDs of the ‘windows’ units through Eq. (9). 2) We masked the window’s region and replaced the activated area (as shown in the third column) with $\min(U_w)$, where U_w is the units for windows. 3) Similar to step-1), we located the ‘cabinets’ units’

channel ids. 4) We obtained the manipulated units by filling $\max(U_c)$ (U_c is the units for the cabinets) on the ‘cabinets’ units’ masked areas. 5) Conditioned by the edited units, the decoder generates new output, in which the area of windows is replaced with cabinets.

We find the proposed framework can also be used to remove artifacts. As shown in the last row of Fig. 5. By manually masking the artifacts area, the ‘artifacts’ units can be found through Eq. (9). Ablating such units improves the quality of translated images.

Table 1: Quantitative Results. We use FID (the lower the better) to measure the quality and LPIPS distance (the higher the better) for the diversity of the translated images.

Tasks	DCLGAN [10]		DSMAP [9]		MUNIT [16]		DRIT [25]		DMIT [56]		Ours	
	FID	LPIPS	FID	LPIPS	FID	LPIPS	FID	LPIPS	FID	LPIPS	FID	LPIPS
Cat \leftrightarrow Human	83.00	0.331	103.65	0.315	89.01	0.362	123.26	0.216	225.49	0.285	72.53	0.469
	77.22	0.501	160.10	0.417	90.50	0.315	166.18	0.259	228.28	0.240	68.66	0.577
Man \leftrightarrow Woman	59.65	0.369	57.34	0.369	66.61	0.264	65.86	0.152	71.23	0.311	57.17	0.398
	54.70	0.394	64.44	0.346	48.41	0.328	98.07	0.161	94.98	0.319	65.12	0.419
Colored-MNIST	59.99	0.433	53.56	0.421	56.37	0.336	69.90	0.403	65.50	0.460	58.97	0.445
	31.06	0.216	35.29	0.401	30.25	0.205	70.53	0.265	37.60	0.180	28.63	0.269
CG \leftrightarrow Real	75.95	0.391	83.31	0.299	74.61	0.340	79.98	0.208	73.11	0.243	70.07	0.503
	70.07	0.476	76.07	0.515	79.53	0.259	85.47	0.211	72.79	0.227	68.71	0.449
Average	63.95	0.389	79.22	0.385	66.91	0.301	94.91	0.247	108.63	0.283	61.27	0.441

4.3 Performance comparison

We qualitatively and quantitatively compared our method with the four state-of-the-art UIT methods: DCLGAN [10], DSMAP [9], MUNIT [16], DRIT [25] and DMIT [56]. We train these four baselines on our datasets with their publicly available implementations. As shown in Table 1, the proposed SCS-UIT framework outperforms the other methods on four datasets in terms of visual quality (*i.e.*, FID) and diversity (*i.e.*, LPIPS). In the first two rows of Fig. 6, we provide a visual comparison on Cat \leftrightarrow Human task. The results show that the geometry of the input image and translated image are more consistent than other methods. The last two rows of Fig. 6 compare the performance of different methods on CG \leftrightarrow Real dataset. Visual results show that our method generates more realistic textures (first row), diverse illuminations (last row). More visual results on these two tasks, Colored-MNIST and Man \leftrightarrow Woman are provided in the supplementary materials.

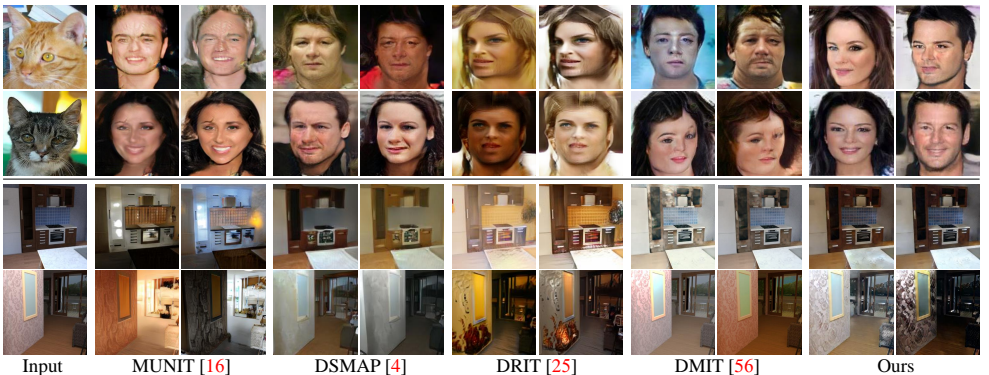


Figure 6: Qualitative comparison on Cat to Human (Top) and on the CG to Real (Bottom).

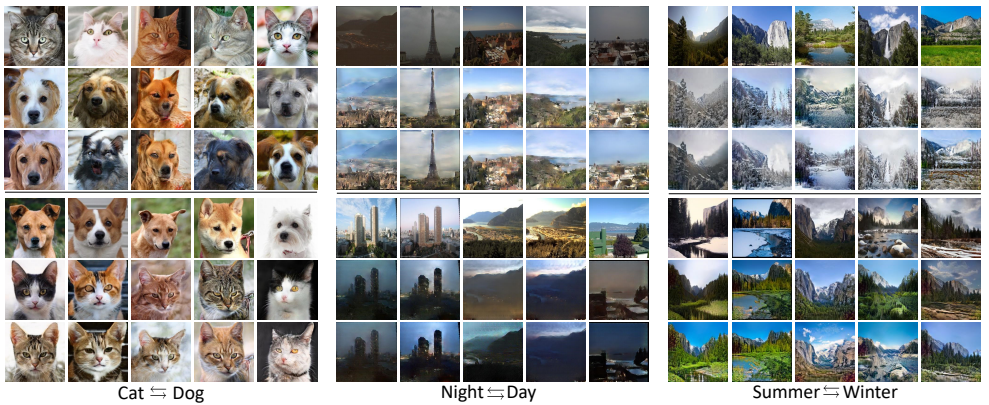


Figure 7: Visual results of our method on other three UIT tasks. For each task, we provide bidirectional translations, which are separated by a solid line. In each translation, we provide 5 inputs (first row) and the corresponding multimodal results (the second and third row).

4.4 Ablation study

To verify the effectiveness of NAIN and SSE, we perform rigorous ablation studies by changing the different parts of our method. We analyze the SSE module by 1) replacing the proposed encoder with two encoders for content code and style code, respectively, and the framework is similar to MUNIT [46]; 2) removing the SSE module (w/o SSE), and extract the content code and style code statically by setting the first quarter of the units as content, and the remaining part as style; 3) replacing NAIN with the original AdaIN (w/o NAIN). Fig. 8 and Table 2 consistently show the final version of SCS-UIT get the best performance both qualitatively and quantitatively. Our full model can improve the content consistency (vs. two encoders), improve the resulting quality (vs. w/o SSE) and remove the water-drop artifact effectively (vs. w/o NAIN).




Figure 8: Visual comparisons of ablation study.

Table 2: Numerical results of ablation study.

	FID ↓	LPIPS ↑
Two Encoders	119.87	0.339
w/o SSE	121.31	0.284
w/o NAIN	78.92	0.402
Ours	72.53	0.449

5 Conclusion

In this paper, we proposed a SCS-UIT framework for unsupervised image-to-image translation. We introduced the domain-invariant high-level vision task for human-perceptive translation. We used an SSE module to separate content code and style code on the basis of relevance between the features and the DI-HV within one unified encoder. The proposed method enables us to interpret and manipulate the translated image. Experimental results on four translation data sets demonstrate that our method outperforms the state-of-the-art methods in terms of visual quality and diversity. This work opens the pathway towards interpreting and manipulating the results in image translation tasks.

References

- [1] Hyojin Bahng, Seungjoo Yoo, Wonwoong Cho, David Keetae Park, Ziming Wu, Xiaojuan Ma, and Jaegul Choo. Coloring with words: Guiding image colorization through text-based palette generation. In *ECCV*, 2018.
- [2] David Bau, Jun-Yan Zhu, Hendrik Strobelt, Bolei Zhou, Joshua B Tenenbaum, William T Freeman, and Antonio Torralba. Gan dissection: Visualizing and understanding generative adversarial networks. In *ICLR*, 2019.
- [3] Sai Bi, Kalyan Sunkavalli, Federico Perazzi, Eli Shechtman, Vladimir G Kim, and Ravi Ramamoorthi. Deep cg2real: Synthetic-to-real translation via image disentanglement. In *ICCV*, 2019.
- [4] Hsin-Yu Chang, Zhixiang Wang, and Yung-Yu Chuang. Domain-specific mappings for generative adversarial style transfers. In *ECCV*, 2020.
- [5] Tian Qi Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. In *NeurIPS*, 2018.
- [6] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *TPAMI*, 2015.
- [7] Hajar Emami, Majid Moradi Aliabadi, Ming Dong, and Ratna Chinnam. Spa-gan: Spatial attention gan for image-to-image translation. *TMM*, 2020.
- [8] Aviv Gabbay and Yedid Hoshen. Improving style-content disentanglement in image-to-image translation. *arXiv preprint arXiv:2007.04964*, 2020.
- [9] Abel Gonzalez-Garcia, Joost Van De Weijer, and Yoshua Bengio. Image-to-image translation for cross-domain disentanglement. In *NeurIPS*, 2018.
- [10] Junlin Han, Mehrdad Shoeiby, Lars Petersson, and Mohammad Ali Armin. Dual contrastive learning for unsupervised image-to-image translation. In *CVPR Workshop*, 2021.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [12] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017.
- [13] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, 2017.
- [14] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018.
- [15] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017.
- [16] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *ECCV*, 2018.

- [17] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017.
- [18] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*. Springer, 2016.
- [19] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019.
- [20] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. *arXiv preprint arXiv:1912.04958*, 2019.
- [21] Junho Kim, Minjae Kim, Hyeonwoo Kang, and Kwanghee Lee. U-gat-it: unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. In *ICLR*, 2020.
- [22] Taeksoo Kim, Moonsu Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. Learning to discover cross-domain relations with generative adversarial networks. In *ICML*, 2017.
- [23] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [24] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *ICCV*, 2017.
- [25] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *ECCV*, pages 35–51, 2018.
- [26] Xueting Li, Sifei Liu, Jan Kautz, and Ming-Hsuan Yang. Learning linear transformations for fast arbitrary style transfer. In *CVPR*, 2019.
- [27] Jianxin Lin, Yingce Xia, Tao Qin, Zhibo Chen, and Tie-Yan Liu. Conditional image-to-image translation. In *CVPR*, 2018.
- [28] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NeurIPS*, 2017.
- [29] Yunfei Liu, Shaodi You, Yu Li, and Feng Lu. Unsupervised learning for intrinsic image decomposition from a single image. In *CVPR*, 2020.
- [30] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Large-scale celebfaces attributes (celeba) dataset. *Retrieved August*, 2018.
- [31] Feifan Lv, Feng Lu, Jianhua Wu, and Chongsoon Lim. Mblen: Low-light image/video enhancement using cnns. In *BMVC*, 2018.
- [32] Feifan Lv, Yinqiang Zheng, Yicheng Li, and Feng Lu. An integrated enhancement solution for 24-hour colorful imaging. In *AAAI*, 2020.

- [33] Liqian Ma, Xu Jia, Stamatios Georgoulis, Tinne Tuytelaars, and Luc Van Gool. Exemplar guided unsupervised image-to-image translation with semantic consistency. *NeurIPS*, 2018.
- [34] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, 2017.
- [35] Youssef Alami Mejjati, Christian Richardt, James Tompkin, Darren Cosker, and Kwang In Kim. Unsupervised attention-guided image-to-image translation. In *NeurIPS*, 2018.
- [36] Ari S Morcos, David GT Barrett, Neil C Rabinowitz, and Matthew Botvinick. On the importance of single directions for generalization. *arXiv preprint arXiv:1803.06959*, 2018.
- [37] Sanghyeon Na, Seungjoo Yoo, and Jaegul Choo. Miso: Mutual information loss with stochastic style representations for multimodal image-to-image translation. *arXiv preprint arXiv:1902.03938*, 2019.
- [38] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016.
- [39] Zhuwei Qin, Fuxun Yu, Chenchen Liu, and Xiang Chen. How convolutional neural network see the world-a survey of convolutional neural network visualization methods. *arXiv preprint arXiv:1804.11191*, 2018.
- [40] Andrés Romero, Pablo Arbeláez, Luc Van Gool, and Radu Timofte. Smit: Stochastic multi-label image-to-image translation. In *ICCV Workshop*, 2019.
- [41] Pravakar Roy, Nicolai Häni, and Volkan Isler. Semantics-aware image to image translation and domain transfer. *arXiv preprint arXiv:1904.02203*, 2019.
- [42] Subhankar Roy, Aliaksandr Siarohin, Enver Sangineto, Nicu Sebe, and Elisa Ricci. Trigan: Image-to-image translation for multi-source domain adaptation. *arXiv preprint arXiv:2004.08769*, 2020.
- [43] Zhiqiang Shen, Mingyang Huang, Jianping Shi, Xiangyang Xue, and Thomas S Huang. Towards instance-level image-to-image translation. In *CVPR*, 2019.
- [44] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [45] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *PRML*, 2017.
- [46] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016.
- [47] Hao Tang, Dan Xu, Nicu Sebe, Yanzhi Wang, Jason J Corso, and Yan Yan. Multi-channel attention selection gan with cascaded semantic guidance for cross-view image translation. In *CVPR*, 2019.
- [48] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne, 2008.

- [49] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Minghui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *TPAMI*, 2019.
- [50] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *ECCV*, 2018.
- [51] Yaxing Wang, Abel Gonzalez-Garcia, Joost van de Weijer, and Luis Herranz. Controlling biases and diversity in diverse image-to-image translation. *arXiv preprint arXiv:1907.09754*, 2019.
- [52] Wayne Wu, Kaidi Cao, Cheng Li, Chen Qian, and Chen Change Loy. Transgaga: Geometry-aware unsupervised image-to-image translation. In *CVPR*, 2019.
- [53] Chao Yang, Taehwan Kim, Ruizhe Wang, Hao Peng, and C-C Jay Kuo. Show, attend, and translate: Unsupervised image translation with self-regularization and attention. *TIP*, 2019.
- [54] Zili Yi, Hao Zhang, Ping Tan, and Minglun Gong. Dualgan: Unsupervised dual learning for image-to-image translation. In *ICCV*, 2017.
- [55] Seungjoo Yoo, Hyojin Bahng, Sunghyo Chung, Junsoo Lee, Jaehyuk Chang, and Jaegul Choo. Coloring with limited data: Few-shot colorization via memory augmented networks. In *CVPR*, 2019.
- [56] Xiaoming Yu, Yuanqi Chen, Shan Liu, Thomas Li, and Ge Li. Multi-mapping image-to-image translation via learning disentanglement. In *NeurIPS*, 2019.
- [57] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *ECCV*. Springer, 2014.
- [58] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [59] Ziqiang Zheng, Zhibin Yu, Haiyong Zheng, Yang Wu, Bing Zheng, and Ping Lin. Generative adversarial network with multi-branch discriminator for cross-species image-to-image translation. *arXiv preprint arXiv:1901.10895*, 2019.
- [60] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017.
- [61] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.