

Marginal Contrastive Correspondence for Guided Image Generation

Fangneng Zhan¹ Yingchen Yu² Rongliang Wu² Jiahui Zhang² Shijian Lu^{*2} Changgong Zhang³
¹ S-Lab, Nanyang Technological University ² Nanyang Technological University ³ Amazon

Abstract

Exemplar-based image translation establishes dense correspondences between a conditional input and an exemplar (from two different domains) for leveraging detailed exemplar styles to achieve realistic image translation. Existing work builds the cross-domain correspondences implicitly by minimizing feature-wise distances across the two domains. Without explicit exploitation of domain-invariant features, this approach may not reduce the domain gap effectively which often leads to sub-optimal correspondences and image translation. We design a Marginal Contrastive Learning Network (MCL-Net) that explores contrastive learning to learn domain-invariant features for realistic exemplar-based image translation. Specifically, we design an innovative marginal contrastive loss that guides to establish dense correspondences explicitly. Nevertheless, building correspondence with domain-invariant semantics alone may impair the texture patterns and lead to degraded texture generation. We thus design a Self-Correlation Map (SCM) that incorporates scene structures as auxiliary information which improves the built correspondences substantially. Quantitative and qualitative experiments on multifarious image translation tasks show that the proposed method outperforms the state-of-the-art consistently.

1. Introduction

Image-to-image translation refers to the image generation conditioned on the certain inputs from other domains [25, 29, 32], and has delivered impressive advances with the emergence of Generative Adversarial Networks (GANs) [5] in recent years. As a typical ill-posed problem, diverse solutions are naturally allowed in image translation tasks as one conditional input could corresponds to multiple image instances. Faithful control of the generation style not only enables diverse generation given certain conditions but also flexible user control of the desired generation. However, yielding high-fidelity images with controllable styles still remains a grand challenge.

*Corresponding author

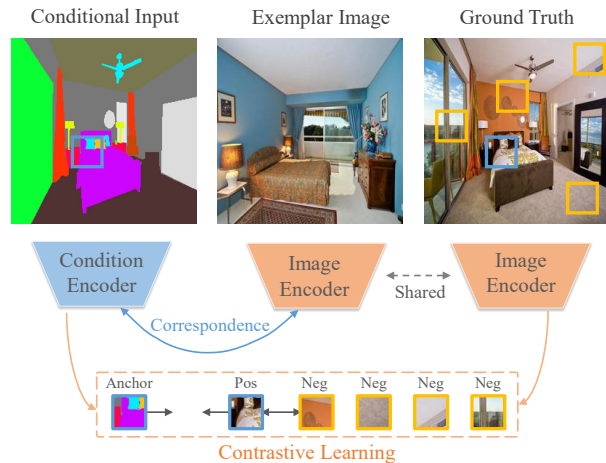


Figure 1. Learning domain-invariant features for building correspondence across domains: We exploit contrastive learning between the *Conditional Input* and the *Ground Truth* by pulling features at the same position closer while pushing those at different positions apart. With the learned *Condition Encoder* and *Image Encoder*, explicit feature correspondences can be built up between the *Conditional Input* and the *Exemplar Image*.

To tackle the style control challenge, one predominant approach employs variational autoencoders (VAEs) to regularize the image style into a Gaussian distribution as in SPADE [25], yet this approach often suffer from compromised generation quality due to the *posterior collapse* phenomenon [18] of VAEs. Another direction is to leverage the styles extracted from the exemplars for guiding the generation. For instance, Zhu *et al.* [47] propose to achieve style control for each semantic region with a dedicated style code. However, the latent codes are typically extracted via a style encoder, which often reflects the global styles only and struggles to capture the detailed structures. Very recently, dense correspondence has been actively explored in exemplar-based image translation, and emerges as a promising approach to achieve faithful style control. Specifically, Zhang *et al.* [40] propose to establish dense semantic correspondence between conditioned input and a given exemplar so as to offer dense style guidance in translation. Zheng *et al.* [43] and CoCosNet v2 [45] enables the establishment of

dense correspondence in high-resolution features via PatchMatch to further enhance the preservation of detailed styles. Zhan *et al.* [34] introduce a general image translation framework that incorporates optimal transport for feature alignment between conditional inputs and style exemplars. As a matching problem across very different domains, the key of building their correspondence lies in how to effectively learn the domain-invariant features to facilitate the proper matching. However, the aforementioned methods implicitly align the domain gaps for correspondence building by minimizing a pseudo pair loss and a L1 loss between features of condition and real image, which does not explicitly model the domain invariant features and may lead to sub-optimal feature correspondence.

In this paper, we propose **Marginal Contrastive Learning Network (MCL-Net)** for exemplar-based image translation, which introduce contrastive learning [2, 6, 31] to effectively extract domain-invariant features for building cross-domain correspondence as shown in Fig. 1. In particular, contrastive learning is applied to the features of the conditional input and its ground truth that are extracted by separate encoders. Each input feature vector is treated as an *anchor*, while the feature vector in the same spatial location is treated as *positive sample* and the remaining feature vectors as *negative samples*. By maximizing the mutual information between condition features and image features, contrastive learning could explicitly yield the domain-invariant features. Furthermore, a marginal contrastive loss (MCL) is proposed to enhance the discriminability of domain-invariant features, which could effectively curb over-smoothed or inaccurate correspondence. A deviation angle in MCL serves as a penalty to the positive anchor for margin preserving.

On the other hand, most previous methods build the correspondence relying on the corresponding local features without knowing the scene structures. It means the scene structure could be impaired while building correspondence, which leads to distorted texture patterns and degraded generation performance. Therefore, the scene structures such as the object shapes should be leveraged to facilitate the correspondence building, especially the preservation of fine texture patterns. Inspired by the observation that the self-attention maps could encode intact image structures even with appearance variations [42], we design a self-correlation map (SCM) mechanism to explicitly represent the scene structures associated with feature and effectively facilitate the correspondence building.

The contributions of this work can be summarized in three aspects. First, we introduce contrastive learning into the exemplar-based image translation framework to explicitly learn domain-invariant features for building correspondence. Second, we propose a novel marginal contrastive loss to boost the feature discriminability in the representa-

tion space, which greatly benefits the establishment of explicit and accurate correspondence. Third, we design a self-correlation map to properly represent the scene structures and effectively facilitate the preservation of texture patterns while building correspondence.

2. Related Work

2.1. Image-to-Image Translation

Due to the superior generation capability, GAN-based image-to-image translation has been extensively investigated and achieved remarkable progress on translating different conditions such as semantic segmentation [10, 25, 32, 37, 39], key points [20, 22, 36, 38] and edge maps [15, 35, 46].

As an ill-posed problem, image translation allows diverse contents or styles generation as long as they comply with the input conditions. However, the styles of generated images are typically determined by the learnt priors of large-scale datasets. To flexibly manipulate the styles for the diverse yet realistic image generation, optimal style control has attracted increasing attention recently. For instance, Huang *et al.* [9] and Ma *et al.* [19] extract and leverage style codes form the exemplar images to guide the generation process via adaptive instance normalization (AdaIN) [8]. Park *et al.* [25] offer a solution that unitizes variational autoencoder (VAE) [13] to encode exemplars for image translation. A style encoder is introduced in Choi *et al.* [3] to ensure the style consistency between the exemplars and the translated images. To specifically control the style in each semantic region, [47] propose a semantic region-adaptive normalization (SEAN). Instead of leveraging latent codes for global style control, Zhang *et al.* [40] propose a framework to enable finer style control by learning dense correspondence between conditional inputs and exemplars. On top of it, Zhang *et al.* [45] introduce PatchMatch to facilitate the learning of dense correspondence on higher resolution, which further preserve better details or styles from the exemplars. However, most existing exemplar-based image translation methods implicitly learn the correspondence by directly apply feature-wised L1 loss to align the domain gaps. This strategy may lead to sub-optimal correspondence as the domain-invariant features are not explicitly explored. We propose to leverage the contrastive learning to explicitly learn domain-invariant features, and boost the discriminability of features via a novel marginal contrastive loss for building more explicit and accurate correspondence.

2.2. Contrastive Learning

Recently, contrastive learning has shown its effectiveness in various computer vision tasks, especially in unsupervised representation learning [2, 6, 23, 26, 31]. Its main idea is to learn a representation by pulling positive samples close to the anchor and pushing negative samples away. Dif-

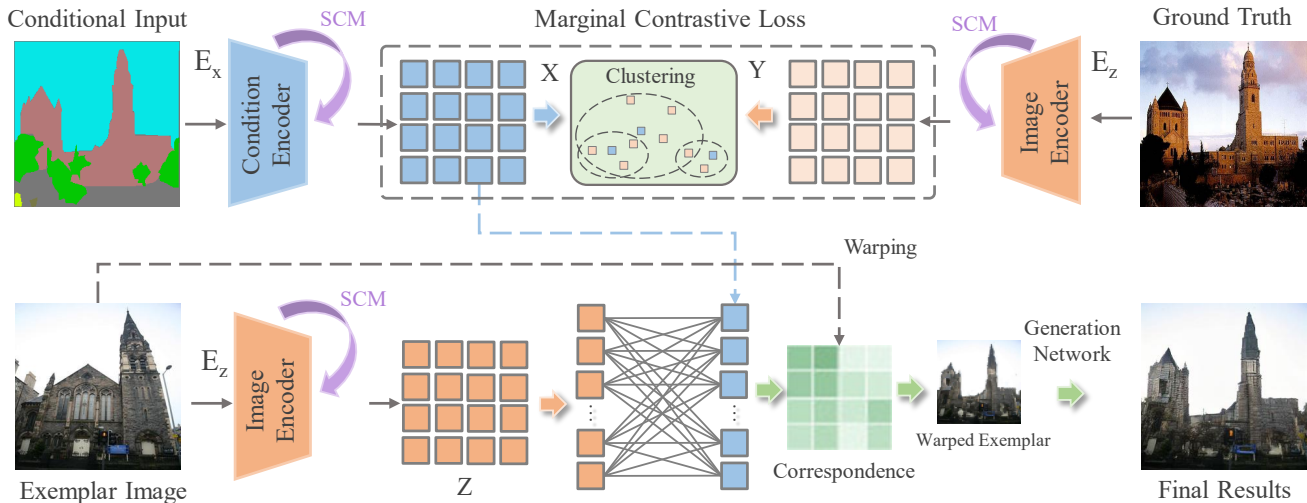


Figure 2. The framework of our proposed conditional image translation network: The *Conditional Input* and *Ground Truth* are fed to feature encoders E_X and E_Z to extract feature vectors X and Y . The proposed self-correlation map (SCM) then encodes structure information for building correspondence, where the proposed *Marginal Contrastive Loss* drives the encoders to learn domain-invariant features. With the shared feature encoder E_Z , domain-invariant features can be extracted from the *Exemplar Image* and feature correspondences between the conditional input and exemplar can be established. The exemplar image can then be warped to align with the conditional input, which provides accurate style guidance for the *Generation Network*.

ferent sampling strategies and contrastive losses have been extensively explored in various downstream tasks. For example, Chen *et al.* [2] and He *et al.* [6] obtain the positive samples by augmenting the original data. Tian *et al.* [31] treat multiple views of the same sample as positive pairs. On top of InfoNCE [23], Park *et al.* [24] introduce PatchNCE to adopt contrastive learning in unpaired image-to-image translation, which treat image patches as samples. In this work, we design a novel marginal contrastive loss based on InfoNCE, which could effectively enhance the discriminability of learned domain-invariant features to help build explicit and accurate correspondence.

3. Proposed Method

Our MCL-Net consists of a correspondence network and a generation network as shown in Fig. 2. The correspondence network employs the proposed marginal contrastive learning to extract domain-invariant features for building correspondence, which is further utilized to warp the exemplars to be aligned with conditional inputs. Under guidance of the warped exemplar and the conditional input, the generation network produces the final translation results. More details to be described in the following subsections.

3.1. Correspondence Network

The correspondence network employs two encoders E_X and E_Z to encode the features of conditional inputs and exemplar images & ground truth, respectively. With

the encoded features $X = [x_1, x_2, \dots, x_N]$ and $Z = [z_1, z_2, \dots, z_N]$ of the conditional input and exemplar image, the correspondence between them can be established by computing the feature-wise Cosine similarity.

To build the feature correspondence effectively between cross-domain images, domain-invariant features is expected to be extracted by the encoders. Previous methods [40, 45] mainly employ a pseudo pairs loss and a L1 loss between conditional inputs and exemplars to drive the learning of feature encoders. However, the pseudo pairs loss does not explicitly encourage the learning of domain-invariant features, while the L1 loss simply minimizes the distance between features in the same spatial location and ignores the distance between features in different spatial locations. We thus introduce contrastive learning to learn domain-invariant features explicitly for building correspondence.

3.2. Contrastive Correspondence

Contrastive learning serves as a powerful tool for unsupervised representation learning by pulling positive samples closer and pushing negative samples apart. To construct the positive and negative pairs for contrastive learning, we feed the ground truth image into encoder E_Z to obtain encoded features $Y = [y_1, y_2, \dots, y_N]$. With a feature x_i from the conditional input as the anchor, the corresponding feature y_i from the ground truth serves as the positive samples, and the remaining $N - 1$ features from the ground truth serve as the negative samples. We normalize vectors onto a unit sphere to prevent the space from collapsing or expanding. Then a




Figure 3. Comparison of warped exemplars and translation results by different methods: Warping and Result denote the warped exemplars and the final translation results. MCL and SCM denote our proposed marginal contrastive learning and self-correlation map, respectively.




Figure 4. Toy illustration of learned feature representations with vanilla contrastive learning and our proposed marginal contrastive learning. Dots indicate image features, angles θ_{ii} and m denote the compactness of the learnt feature clusters and the angular margin penalty, respectively. Zoom-in for details.

noise contrastive estimation [23] framework is employed to conduct contrastive learning between the conditional input and ground truth as below:

$$\mathcal{L}_{xy} = - \sum_{i=1}^N \log \frac{\exp(\frac{x_i \cdot y_i}{\tau})}{\exp(\frac{x_i \cdot y_i}{\tau}) + \sum_{j \neq i}^N \exp(\frac{x_i \cdot y_j}{\tau})} \quad (1)$$

where τ denotes the temperature parameter, ‘ \cdot ’ stands for dot product of vectors. Then a bidirectional contrastive loss \mathcal{L}_{cl} can be formulated by: $\mathcal{L}_{cl} = \mathcal{L}_{xy} + \mathcal{L}_{yx}$.

Marginal Contrastive Learning The contrastive learning enables to encode the domain-invariant features, while building accurate correspondence between two sets of features requires a high discriminability between features. The original contrastive learning tends to produce smooth transition between different feature clusters as given in 1, which may incur smoothed and inaccurate correspondence. Inspired by the additive angular margin loss in ArcFace [4], we propose a marginal contrastive loss to enlarge the separability of features on a hyper-sphere, which yields more explicit and accurate correspondence as shown in Fig. 3.

Algorithm 1 The Pseudo-code of the proposed marginal contrastive loss (MCL).

Input: Conditional input features X, ground truth features Y.

1. $[x_1, x_2, \dots, x_N] = \text{L2Norm}(X)$
2. $[y_1, y_2, \dots, y_N] = \text{L2Norm}(Y)$
3. $\cos \theta_{ii} = x_i \cdot z_i$
4. $\theta_{ii} = \arccos(\cos \theta_{ii})$
5. add margin: $\cos(\theta_{ii} + m)$
6. marginal contrastive loss (MCL):

$$- \sum_{i=1}^N \log \left[\frac{\exp(\cos(\theta_{ii} + m))}{\exp(\cos(\theta_{ii} + m)) + \sum_{j \neq i}^N \exp(\cos \theta_{ij})} \right]$$

Output: MCL.

Specially, as the features for building correspondence are normalized, the loss term \mathcal{L}_{xy} in Eq. (1) can be rewritten as:

$$\mathcal{L}_{xy} = - \sum_{i=1}^N \log \frac{\exp(s \cos \theta_{ii})}{\exp(s \cos \theta_{ii}) + \sum_{j \neq i}^N \exp(s \cos \theta_{ij})}$$

Thus, the embedded features are distributed around each feature center on the hyper-sphere with a radius of s . To simultaneously enhance the intra-class compactness and inter-class discrepancy, we add an additive angular margin penalty m ($m=0.4$ by default) to the positive samples to form marginal contrastive loss (MCL) as below:

$$- \sum_{i=1}^N \log \frac{\exp(s \cos(\theta_{ii} + m))}{\exp(s \cos(\theta_{ii} + m)) + \sum_{j \neq i}^N \exp(s \cos \theta_{ij})}$$

The proposed MCL impose angular margin penalty in the normalized hyper-sphere which effectively enlarges the feature separability as illustrated in Fig. 4.

3.3. Self-Correlation Map

Previous works [40, 45] build the correspondence relying on the corresponding local features without knowing the scene structure. However, building correspondence purely




Figure 5. Illustration of the learnt self-correlation maps at two different spatial locations of the exemplar image.

relying on local features may impair the texture patterns of images, thus leading to degraded generation performance. Intuitively, besides the local feature, scene structures (e.g., object shapes) can serve as auxiliary information to preserve the texture patterns and facilitate the building of correspondence. We thus introduce a **Self-Correlation Map (SCM)** in the correspondence network to utilize the scene structures for building correspondence.

As all regions within same categories exhibit some form of self-correlation, the self-correlation map of a feature effectively encodes the intact object shape information as shown in Fig. 5. By estimating these co-occurrence features for each spatial location, the scene structure of all locations can be explicitly represented. Thus, given encoded features $X = [x_1, x_2, \dots, x_N]$, the self-correlation map (SCM) of a feature can be denoted by:

$$\text{SCM}_i = (x_i)^T \cdot X, \quad i = 0, 1, \dots, N. \quad (2)$$

Fig. 5 illustrates the sample of self-correlation maps. As shown in the figure, the self-correlation maps accurately encode the scene structures in the corresponding location, e.g., the building and sky.

Unlike the original features that would encode various attributes such as style and texture, the self-correlation map only captures the spatial correlation. To leverage SCM for building correspondence, the SCM (64×64 by default) is flattened into column vector (e.g., 4096), followed by fully-connected layer to reduce the vector dimension to 256. Then the feature vector of 256 dimension is concatenated with the encoded feature in corresponding spatial location to compute Cosine similarity for building correspondence. As the proposed marginal contrastive learning is employed to learn domain-invariant feature for building correspondence, the learning of self-correlation maps is also driven by the designed contrastive loss.

3.4. Generation Network

With the warped exemplar to provide style guidance and the conditional input to provide semantic guidance, the generation network aims to generate high-fidelity images

with faithful style to the exemplar and faithful semantic to the conditional input. Specially, the warped exemplar is concatenated with the conditional input to be injected into the generation network through a spatially adaptive de-normalization [25]. The detailed architecture of the generation network is consistent with SPADE [25]. More details of the network structure are available in the supplementary material.

3.5. Loss Functions

The correspondence network and generation network are optimized jointly to learn the cross-domain correspondence and achieve high-fidelity image generation. The conditional input, ground truth, exemplar are denoted by X , Y and Z , respectively. The condition encoder and image encoder in correspondence network are denoted by E_X and E_Z , and the generator and discriminator are denoted by G and D .

Correspondence Network Besides the proposed marginal contrastive loss \mathcal{L}_{mcl} , several other losses are designed to facilitate the learning of cross-domain correspondence. As driven by the contrastive learning, the two feature encoder E_X and E_Z aim to extract domain invariant features (i.e., semantic features), the encoded features from the conditional input X and the corresponding ground truth Y should be consistent, thus yielding a feature consistency loss as defined below:

$$\mathcal{L}_{fst} = \|E_X(X) - E_Z(Y)\|_1. \quad (3)$$

To preserve image information in the warping process, the original exemplar should be recoverable from the warped exemplar through an inverse warping. Thus, a cycle-consistency loss can be formulated as below:

$$\mathcal{L}_{cyc} = \|T^T \cdot T \cdot Z - Z\|_1 \quad (4)$$

where T is the correspondence matrix. Although it is hard to collect the ground truth of warped exemplars, pseudo exemplar pairs can be obtained by treating an augmented real image Y' as the exemplar [40]. We thus can penalize the difference between the warped exemplar and the augmented real image Y' as below:

$$\mathcal{L}_{pse} = \|T \cdot Z - Y'\|_1 \quad (5)$$

Generation Network To achieve high-fidelity image translation, several losses are adopted in the generation network to work collaboratively. The generated image $G(X, Z)$ should be consistent with the ground truth Y in term of the semantics, thus a perceptual loss \mathcal{L}_{perc} [11] is employed to minimize their semantic discrepancy as below:

$$\mathcal{L}_{perc} = \|\phi_l(G(X, Z)) - \phi_l(Y)\|_1 \quad (6)$$

Methods	ADE20K [44]			DeepFashion [16]			CelebA-HQ [17]		
	FID ↓	SWD ↓	LPIPS ↑	FID ↓	SWD ↓	LPIPS ↑	FID ↓	SWD ↓	LPIPS ↑
Pix2pixHD [32]	81.80	35.70	N/A	25.20	16.40	N/A	42.70	33.30	N/A
SPADE [25]	33.90	19.70	0.344	36.20	27.80	0.231	31.50	26.90	0.187
SelectionGAN [30]	35.10	21.82	0.382	38.31	28.21	0.223	34.67	27.34	0.191
SMIS [48]	42.17	22.67	0.416	22.23	23.73	0.240	23.71	22.23	0.201
SEAN [47]	24.84	10.42	0.499	16.28	17.52	0.251	18.88	19.94	0.203
UNITE [34]	25.15	10.13	0.571	13.08	16.65	0.278	13.15	14.91	0.213
CoCosNet [40]	26.40	10.50	0.560	14.40	17.20	0.272	14.30	15.30	0.208
CoCosNet v2 [45]	25.21	9.940	0.564	12.81	16.53	0.283	12.85	14.62	0.218
MCL-Net	24.75	9.852	0.569	12.89	16.24	0.286	12.52	14.21	0.216

Table 1. Comparing MCL-Net with state-of-the-art image translation methods: The comparisons were performed over three public datasets with three widely used evaluation metrics FID, SWD and LPIPS.

Methods	Style Relevance		Semantic
	Color	Texture	Consistency
SPADE [25]	0.874	0.892	0.856
SMIS [48]	0.887	0.907	0.858
SEAN [47]	0.932	0.926	0.860
UNITE [34]	0.963	0.945	0.869
CoCosNet [40]	0.962	0.941	0.862
CoCosNet v2 [45]	0.970	0.948	0.877
MCL-Net	0.966	0.951	0.881

Table 2. Quantitative evaluation of style relevance (color and texture) and semantic consistency on ADE20K.

where ϕ_l denotes l layer of pre-trained VGG-19 [27] model. On the other hand, the generated image $G(X, Z)$ should be consistent with the exemplar in terms of the style. Thus, a contextual loss as described in [21] is formulated as below:

$$\mathcal{L}_{cxt} = -\log\left(\sum_l \mu_l C X_{ij}(\phi_l(Z), \phi_l(Y))\right) \quad (7)$$

where μ_m balances the terms of different VGG layers.

To generate fine-grained image details, an adversarial loss \mathcal{L}_{adv} is introduced with a discriminator D . Thus, the overall objective function of the model is:

$$\mathcal{L} = \min_{E_X, E_Z, G} \max_D (\lambda_1 \mathcal{L}_{cyc} + \lambda_2 \mathcal{L}_{fcst} + \lambda_3 \mathcal{L}_{mcl} + \lambda_4 \mathcal{L}_{perc} + \lambda_5 \mathcal{L}_{cxt} + \lambda_6 \mathcal{L}_{pse} + \lambda_7 \mathcal{L}_{adv}) \quad (8)$$

where λ is weighting parameter. Detailed ablation study of the loss terms can be found in the supplementary material.

4. Experiments

4.1. Experimental Settings

Datasets. We conduct exemplar-based image translation on three different tasks.

- ADE20k [44] consists of 20,210 training images associated with a 150-class segmentation mask. The image translation is conducted with semantic segmentation as the input on this dataset.
- CelebA-HQ [17] collects 30,000 high quality face images. The experiment is conducted with face edge maps as the conditional inputs. Canny edge detector is employed to retrieve the connected face landmarks as face edges.
- DeepFashion [16] contains 52,712 high-quality persons in fashion clothe. The experiment is conducted with the human key-points as the conditional input. OpenPose [1] is employed to retrieve the human pose key-points.

Implementation Details. The correspondence network and generation network are jointly optimized with learning rates $1e-4$ and $4e-4$ for the generator and discriminator. Adam solver with $\beta_1 = 0$ and $\beta_2 = 0.999$ is adopted for the optimization. All experiments were conducted on 4 32GB Tesla V100 GPUs. The resolution for building correspondence is 64×64 . The resolution of generated images is 256×256 for all translation tasks.

4.2. Quantitative Evaluation

Evaluation Metrics. To evaluate the translation results comprehensively, *Fréchet Inception Score (FID)* [7] and *Sliced Wasserstein distance (SWD)* [12] to evaluate the image perceptual quality; *Learned Perceptual Image Patch Similarity (LPIPS)* [41] is employed to evaluate the diversity of translated images. Besides, L1 distance, peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) [33] are also adopted to evaluate the low-level quality of warped images.

To evaluate the style relevance and semantic consistency of translated images, the metrics described in [40] is employed in this work. Specially, as low-level features (*relu1.2, relu2.2*) usually encode image styles such as color and texture, we thus take the average cosine similarity of these layers in pre-trained VGG network [28] as



Figure 6. Qualitative comparison of MCL-Net and state-of-the-art image translation methods over three different types of conditional inputs including semantic segmentation, key-points, and edge map, respectively, from top to bottom.

the score of style relevance ($relu1_2$ for color, $relu2_2$ for texture). As high-level features ($relu3_2$, $relu4_2$, $relu5_2$) tend to encode semantic features, we take the average cosine similarity of these layers in pre-trained VGG network [28] as the score of semantic consistency.

Experimental Results. As shown in Table 1, we can observe that MCL-Net outperforms nearly all compared methods in terms of image quality as measured by FID and SWD and image diversity as measured by LPIPS. Compared with CoCosNet, MCL-Net achieves better FID and SWD as the designed marginal contrastive loss allows to learn domain invariant features for building correspondence and the self-correlation allows to incorporate structure information for building correspondence.

As an exemplar-based image translation task, the generated image should present semantic consistency with conditional inputs and present styles relevance with exemplars. The evaluation results of semantic consistency and style consistency are tabulated in Table. 2. With our marginal contrastive learning for invariant feature learning and SCM for effective utilization of structure information, MCL-Net achieves the best feature correspondence which leads to best style relevance and semantic consistency.

4.3. Qualitative Evaluation

Fig. 6 shows qualitative comparisons of different methods. It can be observed that MCL-Net achieves the best visual quality with faithful styles as exemplars. Pix2pixHD [32] is unable to capture the style from the exemplars. SPADE [25], SMIS [48] adopt variational auto-encoder

Models	FID ↓	SWD ↓	LPIPS ↑
Baseline	32.02	26.87	0.184
+COR	15.77	16.31	0.199
+COR+MCL	12.75	15.03	0.208
+COR+SCM	13.27	15.24	0.219
+COR+MCL+SCM	12.52	14.21	0.216

Table 3. Ablation studies of our MCL-Net designs over CelebA-HQ [17]: The baseline is SPADE [25] that conduct translation without building correspondence. COR denotes building correspondence between conditional inputs and exemplars. MCL and SCM mean to include the proposed marginal contrastive loss and self-correlation map in building correspondence. Model in the last row is the standard MCL-Net.

(VAE) [14] to encode image styles, while this method is unable to capture detailed style feature such as texture and patterns. Exemplar-based translation methods CoCosNet [40] and CoCosNet v2 [45] build cross-domain correspondence to capture detailed style from the exemplars, but they tend to produce blurred generation results without explicitly learning of domain invariant features and lose texture patterns without leveraging image structure information. Our MCL-Net designs MCL to learn domain invariant features explicitly and introduce SCM to leverage the structure information for building correspondence.

Besides image quality, MCL-Net also demonstrates superior diversity in image translation as illustrated in Fig. 7. We can observe that MCL-Net is capable of synthesizing realistic images with faithful style to the given exemplars.

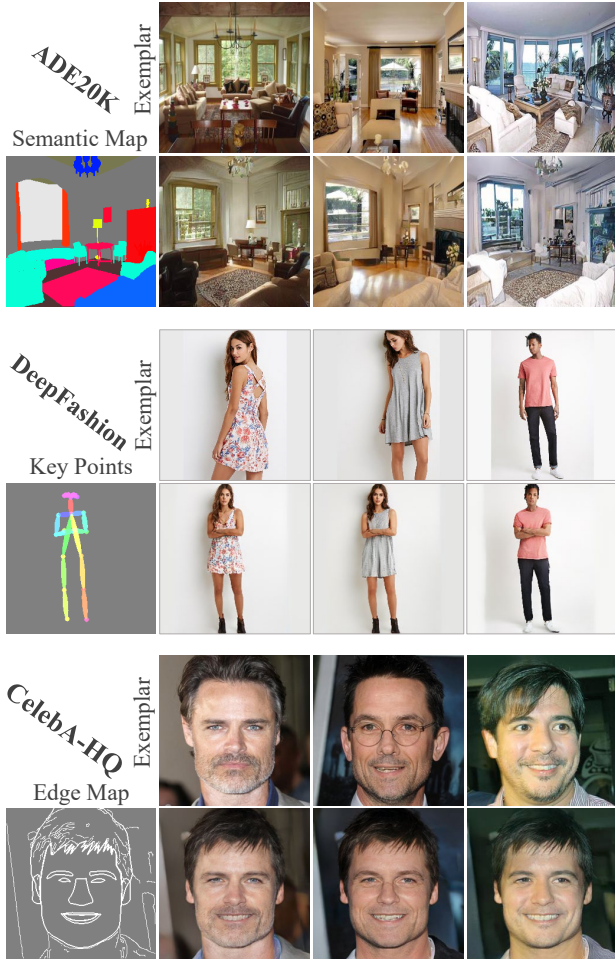


Figure 7. Qualitative illustration of our proposed MCL-Net with different types of conditional inputs and exemplars.

Models	L1 ↓	PSNR ↑	SSIM ↑
Baseline	80.14	28.35	0.766
+MCL ($m=0.1$)	79.32	29.01	0.774
+MCL ($m=0.2$)	78.05	29.44	0.787
+MCL ($m=0.3$)	76.43	29.75	0.795
+MCL ($m=0.4$)	75.27	30.04	0.811
+SCM	76.07	29.67	0.791
+MCL ($m=0.4$)+SCM	73.72	30.94	0.826

Table 4. Parameter studies of marginal contrastive learning (MCL) and self-correlation map (SCM) on warped exemplars on DeepFashion dataset. The baseline is CoCosNet [40], and the last row denotes the standard MCL-Net.

4.4. Discussion

We conduct extensive ablation studies to evaluate our technical designs. Table 3 shows experimental results on

CelebA-HQ. SPADE [25] is selected as the baseline which achieves image translation & manipulation without feature alignment. The performance is clearly improved when feature correspondence (COR) is included to align features (namely CoCosNet [40]). The generation performance can be further improved by including the proposed marginal contrastive learning (MCL) to learn domain invariant features or the self-correlation map (SCM) for building correspondence. Including both MCL and SCM to the model leads to the best generation performance.

To directly evaluate the accuracy of the built correspondence, we evaluate the warped exemplars on DeepFashion dataset where paired data are available by treating person images under different poses as the exemplar and the ground truth. Hence, we can measure the distance between the warped exemplars and the ground truth with L1, PSNR and SSIM [33]. Specially, the margin parameter m plays a key role in MCL, we thus conduct experiments to investigate the effect of m . As shown in Table 4, the vanilla correspondence is selected as the baseline. By varying m from 0.1 to 0.4, we can observe that the correspondence accuracy is improved consistently with the increasing of margin m . Although large margin m contributes to the correspondence accuracy, we find the model tends to be unstable and even fail to converge with a large margin, e.g., $m = 0.5$. We thus select $m = 0.4$ as the default setting in our MCL. Besides, including the designed self-correlation map also improve correspondence accuracy substantially, and combine MCL ($m = 0.4$) and SCM yields the best correspondence accuracy.

5. Conclusions

This paper presents MCL-Net, an exemplar-based image translation framework that introduces a marginal contrastive learning to learn domain invariant feature across conditional input and exemplar for building correspondence. To encode the structure information to facilitate correspondence building, we propose a novel self-correlation map which captures co-occurrence features for each spatial location and represents the image structure explicitly. Quantitative and qualitative experiments show that MCL-Net is capable of generating high-fidelity images with consistent semantic with the conditional input and faithful style to the exemplar.

6. Acknowledgement

This study is supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

- [1] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017. [6](#)
- [2] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. [2](#), [3](#)
- [3] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8188–8197, 2020. [2](#)
- [4] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019. [4](#)
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. [1](#)
- [6] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. [2](#), [3](#)
- [7] 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 *Advances in neural information processing systems*, pages 6626–6637, 2017. [6](#)
- [8] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017. [2](#)
- [9] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 172–189, 2018. [2](#)
- [10] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. [2](#)
- [11] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. [5](#)
- [12] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. [6](#)
- [13] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [2](#)
- [14] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. [7](#)
- [15] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *Proceedings of the European conference on computer vision (ECCV)*, pages 35–51, 2018. [2](#)
- [16] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1096–1104, 2016. [6](#)
- [17] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3730–3738, 2015. [6](#), [7](#)
- [18] James Lucas, George Tucker, Roger Grosse, and Mohammad Norouzi. Don’t blame the elbo! a linear vae perspective on posterior collapse. *arXiv preprint arXiv:1911.02469*, 2019. [1](#)
- [19] Liqian Ma, Xu Jia, Stamatios Georgoulis, Tinne Tuytelaars, and Luc Van Gool. Exemplar guided unsupervised image-to-image translation with semantic consistency. In *International Conference on Learning Representations*, 2018. [2](#)
- [20] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *Advances in neural information processing systems*, pages 406–416, 2017. [2](#)
- [21] Roey Mechrez, Itamar Talmi, and Lihi Zelnik-Manor. The contextual loss for image transformation with non-aligned data. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 768–783, 2018. [6](#)
- [22] Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. Controllable person image synthesis with attribute-decomposed gan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5084–5093, 2020. [2](#)
- [23] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [2](#), [3](#), [4](#)
- [24] Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *European Conference on Computer Vision*, pages 319–345. Springer, 2020. [3](#)
- [25] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2337–2346, 2019. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [26] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, Jasmine Hsu, Eric Jang, Stefan Schaal, Sergey Levine, and Google Brain. Time-contrastive networks: Self-supervised learning from video. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1134–1141. IEEE, 2018. [2](#)
- [27] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [6](#)

- [28] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 6, 7
- [29] Hao Tang, Dan Xu, Gaowen Liu, Wei Wang, Nicu Sebe, and Yan Yan. Cycle in cycle generative adversarial networks for keypoint-guided image generation. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 2052–2060, 2019. 1
- [30] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2417–2426, 2019. 6
- [31] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019. 2, 3
- [32] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018. 1, 2, 6, 7
- [33] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6, 8
- [34] Fangneng Zhan, Yingchen Yu, Kaiwen Cui, Gongjie Zhang, Shijian Lu, Jianxiong Pan, Changgong Zhang, Feiying Ma, Xuansong Xie, and Chunyan Miao. Unbalanced feature transport for exemplar-based image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 2, 6
- [35] Fangneng Zhan, Yingchen Yu, Rongliang Wu, Kaiwen Cui, Aoran Xiao, Shijian Lu, and Ling Shao. Bi-level feature alignment for semantic image translation & manipulation. *arXiv preprint*, 2021. 2
- [36] Fangneng Zhan, Yingchen Yu, Rongliang Wu, Changgong Zhang, Shijian Lu, Ling Shao, Feiying Ma, and Xuansong Xie. Gmlight: Lighting estimation via geometric distribution approximation. *arXiv preprint arXiv:2102.10244*, 2021. 2
- [37] Fangneng Zhan, Yingchen Yu, Rongliang Wu, Jiahui Zhang, and Shijian Lu. Multimodal image synthesis and editing: A survey. *arXiv preprint arXiv:2112.13592*, 2021. 2
- [38] Fangneng Zhan, Changgong Zhang, Yingchen Yu, Yuan Chang, Shijian Lu, Feiying Ma, and Xuansong Xie. Em-light: Lighting estimation via spherical distribution approximation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3287–3295, 2021. 2
- [39] Fangneng Zhan, Jiahui Zhang, Yingchen Yu, Rongliang Wu, and Shijian Lu. Modulated contrast for versatile image synthesis. *arXiv preprint arXiv:2203.09333*, 2022. 2
- [40] Pan Zhang, Bo Zhang, Dong Chen, Lu Yuan, and Fang Wen. Cross-domain correspondence learning for exemplar-based image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5143–5153, 2020. 1, 2, 3, 4, 5, 6, 7, 8
- [41] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6
- [42] Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. The spatially-correlative loss for various image translation tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16407–16417, 2021. 2
- [43] Haitian Zheng, Zhe Lin, Jingwan Lu, Scott Cohen, Jianming Zhang, Ning Xu, and Jiebo Luo. Semantic layout manipulation with high-resolution sparse attention. *arXiv preprint arXiv:2012.07288*, 2020. 1
- [44] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 633–641, 2017. 6
- [45] Xingran Zhou, Bo Zhang, Ting Zhang, Pan Zhang, Jianmin Bao, Dong Chen, Zhongfei Zhang, and Fang Wen. Cocosnet v2: Full-resolution correspondence learning for image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11465–11475, 2021. 1, 2, 3, 4, 6, 7
- [46] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman. Toward multimodal image-to-image translation. In *Advances in neural information processing systems*, pages 465–476, 2017. 2
- [47] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. Sean: Image synthesis with semantic region-adaptive normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5104–5113, 2020. 1, 2, 6
- [48] Zhen Zhu, Zhiliang Xu, Ansheng You, and Xiang Bai. Semantically multi-modal image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5467–5476, 2020. 6, 7