

CFFT-GAN: Cross-Domain Feature Fusion Transformer for Exemplar-Based Image Translation

Tianxiang Ma^{1,2*}, Bingchuan Li³, Wei Liu³, Miao Hua³, Jing Dong^{2†}, Tieniu Tan^{2,4}

¹School of Artificial Intelligence, University of Chinese Academy of Sciences

²CRIPAC & NLPR, Institute of Automation, Chinese Academy of Sciences

³ByteDance Ltd, Beijing, China

⁴Nanjing University

tianxiang.ma@cripac.ia.ac.cn, {libingchuan, huamiao}@bytedance.com, liujikun63@gmail.com, {jdong, tnt}@nlpr.ia.ac.cn

Abstract

Exemplar-based image translation refers to the task of generating images with the desired style, while conditioning on certain input image. Most of the current methods learn the correspondence between two input domains and lack the mining of information within the domains. In this paper, we propose a more general learning approach by considering two domain features as a whole and learning both inter-domain correspondence and intra-domain potential information interactions. Specifically, we propose a Cross-domain Feature Fusion Transformer (CFFT) to learn inter- and intra-domain feature fusion. Based on CFFT, the proposed CFFT-GAN works well on exemplar-based image translation. Moreover, CFFT-GAN is able to decouple and fuse features from multiple domains by cascading CFFT modules. We conduct rich quantitative and qualitative experiments on several image translation tasks, and the results demonstrate the superiority of our approach compared to state-of-the-art methods. Ablation studies show the importance of our proposed CFFT. Application experimental results reflect the potential of our method.

Introduction

Exemplar-based image translation is a subdivision of image translation, which aims to generate a image with exemplar image style, while maintaining the content information of the input image. Such methods have a wide range of promising applications, such as style transfer, face editing, scene transformation, etc.

Previous extensive studies (Ma et al. 2018; Wang et al. 2019; Huang and Belongie 2017; Park et al. 2019; Zhang et al. 2020; Zhou et al. 2021; Zhan et al. 2021) have made great progress in exemplar-based image translation. However, current methods mostly focus on learning the correspondence between the two input domains and ignore the potentially useful information interactions within each domain. In this paper, we propose a more general approach that combines the two domain features of the inputs into a whole, and then uses a transformer-based model to learn

both the inter-domain correspondence and the information interactions within each domain.

Recently, a number of transformer-based models (Dosovitskiy et al. 2021; Carion et al. 2020; Wang et al. 2020; Liu et al. 2021) have been applied to vision tasks. These works outperform CNN-based methods in their respective fields. Nevertheless, training a pure transformer model is time-consuming and highly dependent on the amount of data. Therefore, we utilize a hybrid structure of CNN and transformer to generate better results with limited amount of data. We first use CNN-based encoders to obtain features of the input image and the exemplar image and concatenate them into a whole. Then we propose a novel Cross-domain Feature Fusion Transformer (CFFT) to fuse the whole feature. The CFFT consists of a Feature Fusion Network and a Hierarchical Transformer Network. The former is used for feature fusion and the latter for learning fine-grained spatial alignment at multiple scales. By the improved transformer-based network, CFFT is able to continuously learn self-attention and cross-attention between two input features, so that the model can learn better inter-domain correspondence and potentially beneficial correlation information within each domain.

Based on the CFFT model, we propose CFFT-GAN, a general exemplar-based image translation method. It decouples and fuses features from different domains and leverages generative adversarial network to learn image translation. In particular, we employ encoders to extract features from different image domains and perform feature fusion through CFFT. Then, the fused feature is fed into the spatially-adaptive generator to synthesize the translation results. In addition, our CFFT model is a pluggable design, and we can achieve multi-domain image translation by cascading CFFT modules.

We validate the different types of exemplar-based image translation capabilities of our method on many high-resolution image datasets, including pose-to-image, edge-to-image, sketch-to-image, and mask-to-image. These adequate and diverse experiments have demonstrated the validity of our model and show that our approach has better performance than state-of-the-art methods. Ablation studies and feature visualization results indicate the importance of our proposed CFFT module. Extended multi-domain exper-

* Author did this work during his internship at ByteDance.

† Corresponding author.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

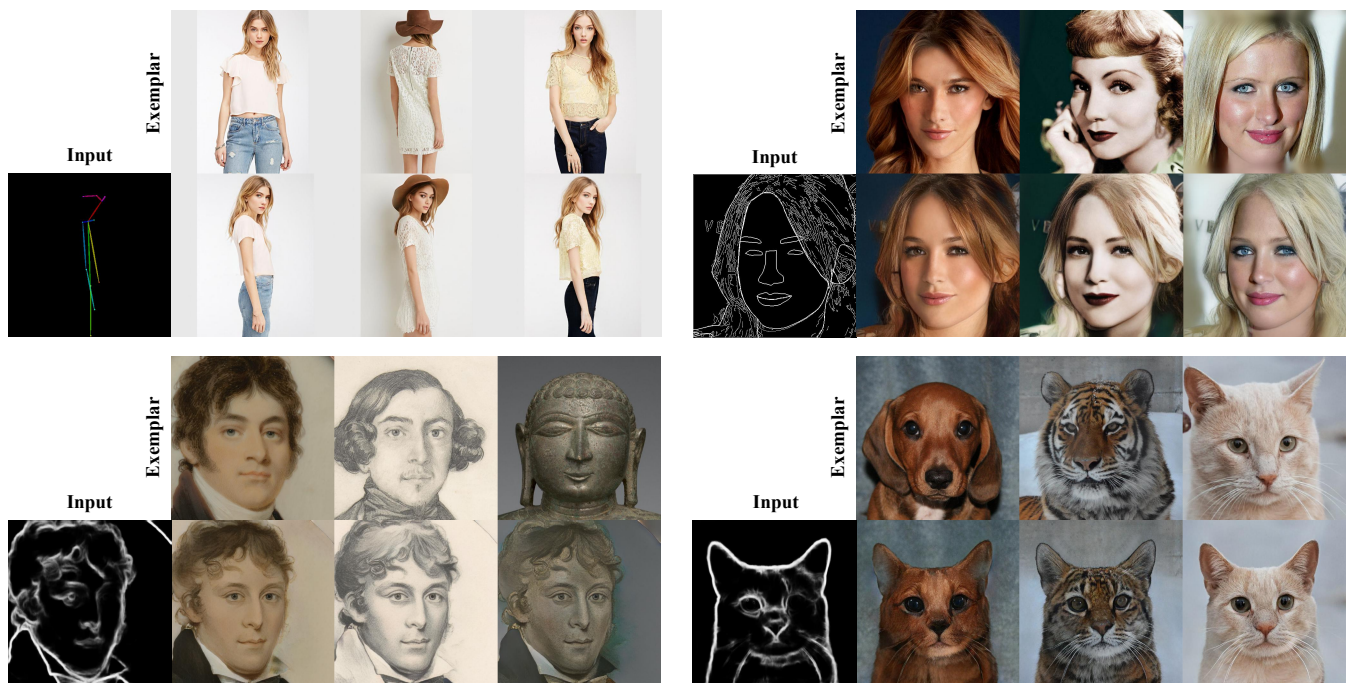


Figure 1: Exemplar-based image translation results of our method on Deepfashion dataset, CelebA-HQ dataset, MetFace dataset, and AFHQ dataset.

iments reflect the potential of our approach.

Our contributions can be summarized as follows:

- We present a general exemplar-based image translation method, CFFT-GAN, which using Cross-domain Feature Fusion Transformer module to learn both intra-domain and inter-domain correlation for more effective image translation.
- We propose a Hierarchical Transformer (Hiformer) network in CFFT for fine-grained spatial alignment at multiple scales.
- Benefiting from the pluggable nature of CFFT, we utilize cascaded CFFT modules to achieve multi-domain image translation.
- Rich validation and comparison experiments show that CFFT-GAN can achieve remarkable results in various image translation tasks and also demonstrate the superiority of our approach over SOTA methods.

Related Work

Image Translation

Image translation aims to learn the transformation between different image domains by training a conditional generative adversarial network. Some methods (Isola et al. 2017; Wang et al. 2018; Park et al. 2019) use paired data for explicit supervision, while others (Zhu et al. 2017a; Choi et al. 2018, 2020; Yi et al. 2017; Park et al. 2020) implement unsupervised training without paired data. Our method works well under both conditions. In addition image translation has been used in a variety of scenarios, such as image editing

(Yu et al. 2018; Wu et al. 2020; Shen et al. 2020), domain adaptation (Liu, Zhang, and Wang 2021; Li et al. 2020), pose transfer (Siarohin et al. 2018; Li, Huang, and Loy 2019; Ren et al. 2020; Ma et al. 2021). Recently, exemplar-based image translation (Ma et al. 2018; Huang et al. 2018; Wang et al. 2019; Saito, Saenko, and Liu 2020; Zhang et al. 2020; Zhou et al. 2021; Zhan et al. 2021) has received more and more attention because of its flexibility and applicability. Usually exemplar-based image translation refers to the transfer of stylistic features from an exemplar image to some type of label image, such as edges (Zhu et al. 2017b; Lee et al. 2018), key points (Ma et al. 2017; Men et al. 2020; Ma et al. 2021), segmentation masks (Wang et al. 2018; Park et al. 2019), etc. Some approaches (Huang et al. 2018; Men et al. 2020; Ma et al. 2021) use AdaIN (Huang and Belongie 2017) to transfer style information from the exemplar image into the label image. More recently, CoCosNet v1 (Zhang et al. 2020) uses correlation matrices to learn cross-domain correspondences. UNITE (Zhan et al. 2021) introduces optimal transport for proper feature alignment and faithful style control in image translation. CoCosNet v2 (Zhou et al. 2021) learn a full-resolution correspondence with a hierarchical GRU-assisted PatchMatch (Barnes et al. 2009). In this paper, we propose a new method using the transformer-based network to learn cross-domain feature fusion.

Transformer

Transformer was first proposed by (Vaswani et al. 2017) for machine translation. Early research on applying transformer to vision tasks has not received extensive attention, while the recent proposal of ViT (Dosovitskiy et al. 2021) model

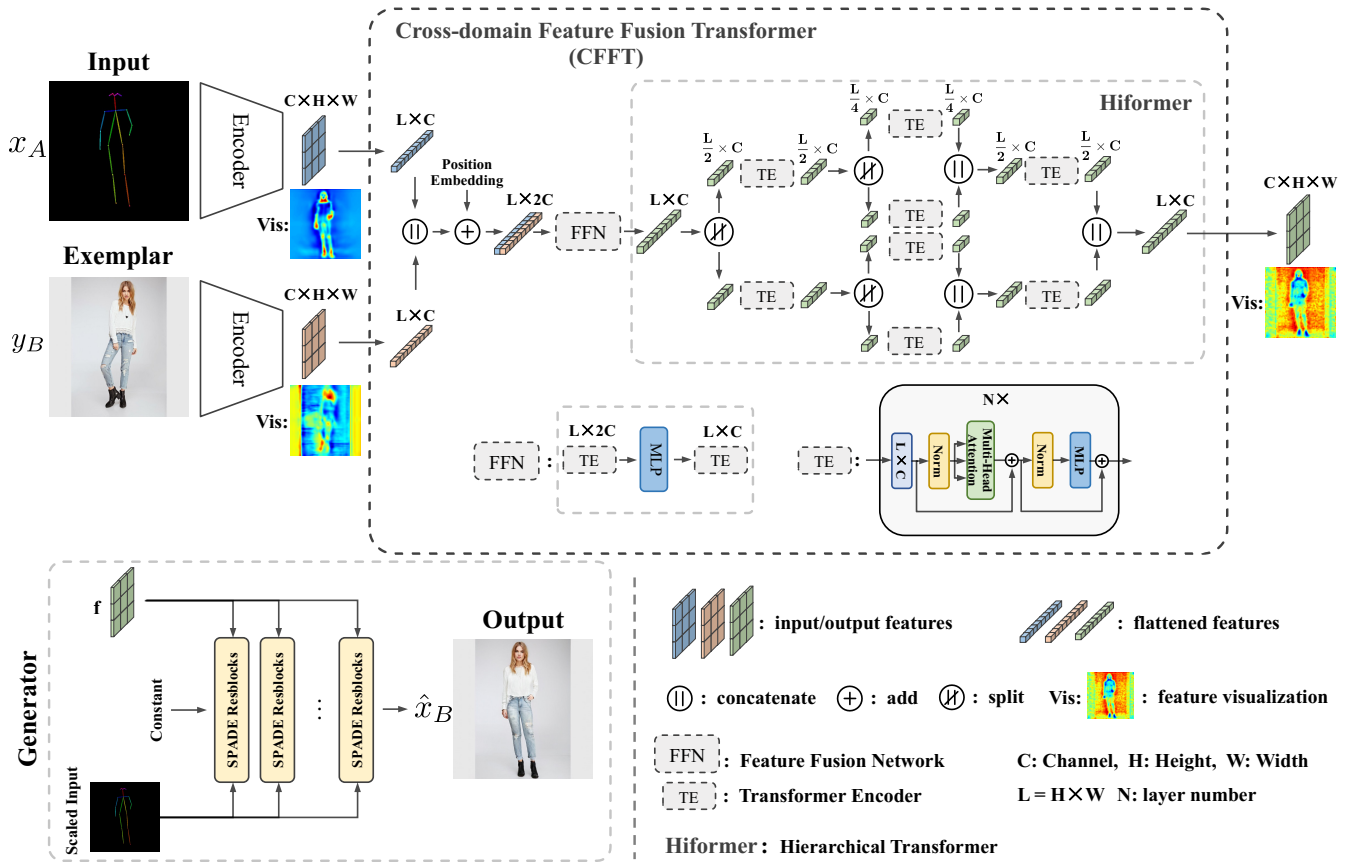


Figure 2: Overall of our CFFFT-GAN model for exemplar-based image translation. Two Encoders are used to extract different domain features. The CFFFT module first concatenates different domain features as a whole feature, and performs feature fusion by FFN, and then learns fine-grained spatial alignment of the feature by Hiformer network. Finally, the output feature from CFFFT is input to the spatially-adaptive generator to synthesize translation result.

has set off a wave of research on visual transformer models. The idea of ViT is very straightforward, which is to transfer the original transformer model as invariant as possible to the image classification task. After ViT, more and more vision tasks started to use the transformer model and its variants. DETR (Carion et al. 2020) applies transformer to object detection task with remarkable results. Axial-deeplab (Wang et al. 2020) applies axial attention to image segmentation. ViViT (Arnab et al. 2021) proposes a video classification model based purely on visual transformer. Swin transformer (Liu et al. 2021) improves the performance of transformer for various visual tasks by further improving the computation of image patch window. Although transformer has achieved good results in various classical vision tasks, there are relatively few studies related to the use of transformer in dense estimation tasks like image generation. TransGAN (Jiang, Chang, and Wang 2021) uses a pure transformer structure to design the generator and discriminator to train the adversarial model, but the resolution of the generated images is small. Recently, ViTGAN (Lee et al. 2021) combines ViT and GAN and designs a new regularization technique to achieve comparable results to CNN-based

GAN models on many standard image generation benchmarks. These studies demonstrate that transformer models can also play an effective role in image generation. However, the mainstream image translation approaches are still using CNN-based networks. Because image translation requires that the generated image has accurate semantic information about the conditional input as well as style information about the exemplar image, which demands learning the correspondence between different image domains. In this paper, we propose a cross-domain feature fusion transformer model for exemplar-based image translation. This model can better learn the correspondences between different domains and the potential connections within each domain.

Approach

Exemplar-based image translation refers to transfer an input image x_A from domain A to domain B , and make it have the stylistic texture of exemplar y_B , while maintaining the content of x_A . Our proposed CFFFT-GAN first transform x_A and y_B into feature space F by two CNN-based encoders, then fuse them using Cross-domain Feature Fusion Transformer, and generate the translation result via a SPADE-based (Park

et al. 2019) residual generator. The overall of our method is illustrated in Figure 2.

Encoders

For exemplar-based image translation, the inputs are two different types of image domains, content input domain A and style exemplar domain B . There is large semantic gaps between them, which is not conducive to the feature fusion later. To learn correspondence better, we need to map the two image domains into the same latent feature space F . Specifically, we let $\mathcal{E}_{A \rightarrow S}$ and $\mathcal{E}_{B \rightarrow S}$ represent the encoders of two domains, then the domain aligned latent features can be expressed as:

$$\begin{aligned} x_F &= \mathcal{E}_{A \rightarrow F}(x_A; \theta_{\mathcal{E}_A}), \\ y_F &= \mathcal{E}_{B \rightarrow F}(y_B; \theta_{\mathcal{E}_B}), \end{aligned} \quad (1)$$

where θ is the learnable parameters for each encoder, and x_F and $y_S \in \mathbb{R}^{C \times H \times W}$ represent the learned domain aligned features, which are used for later feature fusion. The domain alignment of features is trained by Equation 4.

Cross-Domain Feature Fusion Transformer

We propose a novel Cross-domain Feature Fusion Transformer (CFFT) model, shown on the Figure 2. The CFFT is designed to learn not only inter-domain correspondences, but also feature correlations within domains. Therefore, we can more fully utilize the features of both domains for image translation. Cross-domain feature pairs are first flattened into feature vectors. Then we concatenate the flattened features into a whole along the channel dimension. There are a total of $L = H \times W$ tokens, and the dimension of each token is $1 \times 2C$. The purpose of this operation is that the transformer-based network can learn the correlation information within each feature and also indirectly learn the correspondence between different domain features.

Feature Fusion Network. After adding the position embedding, we then use Feature Fusion Network (FFN) to learn feature fusion and semantic coarse matching. FFN consists of two Transformer Encoders (TE) (Vaswani et al. 2017), each TE has a module layer number $N=3$. The two TEs are connected through a MLP layer for compressing feature dimension from $2C$ to C . In this way, FFN achieves a feature fusion role, and based on the transformer’s self-attention mechanism, the module is able to learn both the inter-domain and intra-domain correlations within the total feature.

Hierarchical Transformer. To decouple and match the two features more finely, we propose a Hierarchical Transformer (Hiformer) after FFN. It presents a stepped structure that continuously splits features in spatial dimension and performs self-learning optimization for each spatial part of the feature layer by layer. For each TE in Hiformer, we use a module layer number $N=2$. This approach implements a transformer-based feature learning process from coarse to fine. The number of feature splitting levels is adjustable, and we use 3 levels in this paper. We then gradually fuse the local features by same structure.

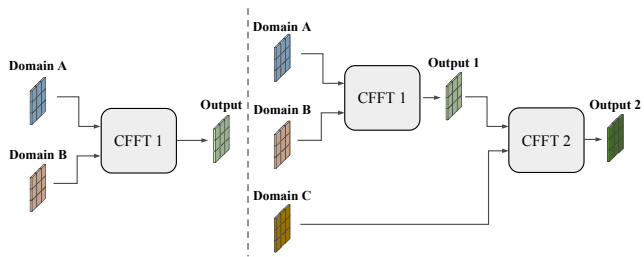


Figure 3: Our method can achieve multi-domain image translation by cascading CFFT modules.

With the above network structure, we can obtain the output feature that have been decoupled and spatially aligned, as shown in the feature visualizations of Figure 2.

Spatially-Adaptive Generator

The role of the generator network is to synthesize \hat{x}_B with the content of input image and the style of exemplar image. In the previous step, we have obtained the fused and aligned feature named \mathbf{f} , we just need to feed it into the generator network. Therefore, we introduce a spatially-adaptive denormalization (SPADE) (Park et al. 2019) structure as the generator. We concatenate the feature \mathbf{f} with the scaled input image into \mathbf{w} , where scaled input image can help the initial training of the model. Then we feed \mathbf{w} into the SPADE-based residual blocks as modulation parameters layer by layer. Specifically, for a particular feature layer $F^i \in \mathbb{R}^{C_i \times H_i \times W_i}$ in the generator, we have the following formula,

$$F_{c,h,w}^i = \gamma_{h,w}^i(\mathbf{w}_i) \times \frac{F_{c,h,w}^i - \mu_{h,w}^i}{\sigma_{h,w}^i} + \beta_{h,w}^i(\mathbf{w}_i), \quad (2)$$

where $\mu_{h,w}^i$ and $\sigma_{h,w}^i$ are the calculated mean and standard deviation across channel dimension, $\gamma_{h,w}^i$ and $\beta_{h,w}^i$ are learnable weight networks, \mathbf{w}_i is a copy of \mathbf{w} . Finally, the generated result of the image translation can be expressed as:

$$\hat{x}_B = \mathcal{G}(\mathbf{w}; \theta_G), \quad (3)$$

where θ_G is the learnable parameters of the generator \mathcal{G} .

Multi-Domain Image Translation

Our proposed CFFT is a pluggable module design. When the inputs have more than two image domains, we can also easily decouple and fuse information from multiple domains by cascading CFFT modules. As shown in schematic Figure 3, we first learn a feature fusion using CFFT1 for domains A and B, then introduce a third domain C and jointly learn a feature fusion for all three domains by CFFT2. This way of learning facilitates the progressive fusion of features, allowing for better correspondence learning between several different domains. As shown in Figure 5, our method can introduce a third domain (pose) on top of the dual-domain image translation to achieve both image translation and face animation.

Loss Functions

We train the CFFT-GAN model end-to-end. Usually, semantically aligned data for different domains is accessible, for example x_A and x_B , whereas triplets data $(x_A, y_B; x_B)$ may not be accessible. We therefore construct the pseudo-tuple $(x_A, \tilde{x}_B; x_B)$, where $\tilde{x}_B = \mathcal{A}(x_B)$, \mathcal{A} denotes a collection of data augmentation operations, such as horizontal flip, geometric deformation, random crop, etc. The following loss functions are used for the training process.

Domain Alignment Loss. We need the features x_F and y_F learned by the encoders to be in the same feature domain, as explained in subsection "Encoders". Therefore, we let different domain features under the same semantic be aligned.

$$\mathcal{L}_{\text{align}} = \|\mathcal{E}_{A \rightarrow S}(x_A) - \mathcal{E}_{B \rightarrow S}(x_B)\|_1. \quad (4)$$

Feature Matching Loss. Following previous work (Johnson, Alahi, and Fei-Fei 2016; Zhang et al. 2020), we used the pre-trained VGG-19 network to extract the features of the generated translation result and the corresponding Ground Truth separately, and let the features at each layer match.

$$\mathcal{L}_{\text{match}} = \sum_l \mu_l \|\phi_l(\hat{x}_B) - \phi_l(x_B)\|_1, \quad (5)$$

where ϕ_l is the output features of layer l in the VGG-19 network, and μ_l is the corresponding weight for each layer.

Translation Loss. We want the model to output translation image that has semantic content of the input image and the stylized texture of the exemplar image. Therefore, We utilize perceptual loss to reduce semantic differences and contextual loss (Mechrez, Talmi, and Zelnik-Manor 2018) to constrain the stylistic consistency. Specifically, the perceptual loss is shown as follows,

$$\mathcal{L}_{\text{perc}} = \|\phi_h(\hat{x}_B) - \phi_h(x_B)\|_2, \quad (6)$$

where ϕ_h represents the output of a high-level layer in the pre-trained VGG-19 network, and this layer contains mainly high-level semantic information about the input image. The contextual loss (CX) is written as follows,

$$\mathcal{L}_{\text{CX}} = \sum_l \omega_l [-\log(\text{CX}(\phi_l(\hat{x}_B), \phi_l(y_B)))] , \quad (7)$$

where CX stands for contextual similarity, ϕ_l is the l th layer of low-level features in the pre-trained VGG-19 network, which contain mainly the rich stylized texture of the image, and ω_l is the weight coefficients of the different layers.

Adversarial Loss. We also use a discriminator network to distinguish between the real samples and the generated samples, and train it together with our image translation model. The adversarial loss function is defined as,

$$\begin{aligned} \mathcal{L}_{\text{adv}}^{\mathcal{D}} &= -\mathbb{E}[h(\mathcal{D}(y_B))] - \mathbb{E}[h(-\mathcal{D}(\mathcal{G}(x_A, y_B)))] , \\ \mathcal{L}_{\text{adv}}^{\mathcal{G}} &= -\mathbb{E}[\mathcal{D}(\mathcal{G}(x_A, y_B))] , \end{aligned} \quad (8)$$

where $h(t) = \min(0, -1 + t)$ represents the hinge function to regularize the discriminator.

Total Loss. In conclusion, we optimize the following overall objective.

$$\begin{aligned} \mathcal{L} &= \lambda_{\text{align}} \mathcal{L}_{\text{align}} + \lambda_{\text{match}} \mathcal{L}_{\text{match}} + \lambda_{\text{perc}} \mathcal{L}_{\text{perc}} \\ &+ \lambda_{\text{CX}} \mathcal{L}_{\text{CX}} + \lambda_{\text{adv}} (\mathcal{L}_{\text{adv}}^{\mathcal{D}} + \mathcal{L}_{\text{adv}}^{\mathcal{G}}) , \end{aligned} \quad (9)$$

where weights λ are used to balance the individual objectives.

Experiments

Datasets. To verify the generality of our approach, we perform experiments on several widely used public datasets.

- Deepfashion (Liu et al. 2016) contains 52,712 human images with various appearances and poses. Human image paired data is used in our experiments. We also utilize Openpose (Cao et al. 2017) to extract the key points of the human image as input pose.
- CelebA-HQ (Liu et al. 2015) is comprised of 30,000 high resolution images of faces in a variety of styles. We connect its face landmarks to form the face edge, and use the Canny edge detection algorithm to get the background edge, and finally combine them as the input.
- MetFaces (Karras et al. 2020) is a high resolution art-style portrait dataset, which contains 1,336 images of human faces. We use HED (Xie and Tu 2015) to extract the sketch of the face images.
- ADE20K (Zhou et al. 2017) contains 25,574 scene images with 150-class semantic segmentation annotation. We use this dataset to validate the image translation from mask to scene.
- AFHQ (Choi et al. 2020) is a animal faces dataset consisting of 15,000 high-quality images. We also use HED to extract the sketch of the animal faces.

For better comparison with the SOTA methods and to demonstrate the good performance of our method, all datasets are experimented with 512 resolution images, except for ADE20K where 256 resolution images are used.

Baselines. We compare our method with SOTA exemplar-based image translation methods. 1) CoCosNet v1 (Zhang et al. 2020) summarized various types of image translation tasks into exemplar-based image translation for the first time. It proposes a cross-domain correspondence learning method, which deforms the exemplar image by calculating the correlation matrix of two domain features. 2) UNITE (Zhan et al. 2021) introduces optimal transport for proper feature alignment and faithful style control in image translation. 3) CoCosNet v2 (Zhou et al. 2021) learn a full-resolution correspondence, and purpose a hierarchical GRU-assisted PatchMatch method for efficient correspondence computation. Because we propose a general image translation method, there is no comparison with task-specific methods. Besides, we do not compare with the approaches that directly learn image translation and fail to use an exemplar.

Evaluation Metrics. We adopt the commonly used image generation evaluation metrics to assess the performance of the various methods. Fréchet Inception Score (FID) (Heusel et al. 2017) is the metric used to evaluate the quality and diversity of generated images. Sliced Wasserstein distance (SWD) (Karras et al. 2017) is applied to measure the statistical distance of low-level patch distribution between synthesized images and real images. Learned Perceptual Image Patch Similarity (LPIPS) is a semantic similarity evaluation metric, and we use it to evaluate the decoupling and semantic alignment of the models on Deepfashion dataset, since only Deepfashion has paired data to compute this metric.

For exemplar-based image translation, the generated image should have the same semantic content as the input and

	Deepfashion			CelebA-HQ		Metfaces		ADE20K	
	FID ↓	SWD ↓	LPIPS ↓	FID ↓	SWD ↓	FID ↓	SWD ↓	FID ↓	SWD ↓
CoCosNet v1	14.4	29.0	0.2386	19.4	22.3	25.6	24.3	26.4	10.5
UNITE	13.1	26.6	0.2371	18.3	21.6	24.1	23.5	25.1	10.1
CoCosNet v2	12.2	24.6	0.2245	16.2	19.7	23.3	22.4	25.2	9.9
Ours	8.8	19.0	0.2070	12.4	18.1	22.6	21.3	25.4	10.1

Table 1: Quantitative comparison results of image quality between our method and SOTA methods.

	Deepfashion	CelebA-HQ	Metfaces	ADE20K
CoCosNet v1	0.924	0.945	0.941	0.862
UNITE	0.944	0.948	0.956	0.839
CoCosNet v2	0.959	0.955	0.963	0.877
Ours	0.975	0.961	0.968	0.878

Table 2: Quantitative evaluation of semantic consistency. The higher the score, the better the evaluation metric.

	Deepfashion	CelebA-HQ	Metfaces	ADE20K
CoCosNet v1	0.959	0.947	0.944	0.952
UNITE	0.966	0.953	0.965	0.955
CoCosNet v2	0.974	0.966	0.964	0.959
Ours	0.982	0.974	0.971	0.954

Table 3: Quantitative evaluation of style similarity. The higher the score, the better the evaluation metric.

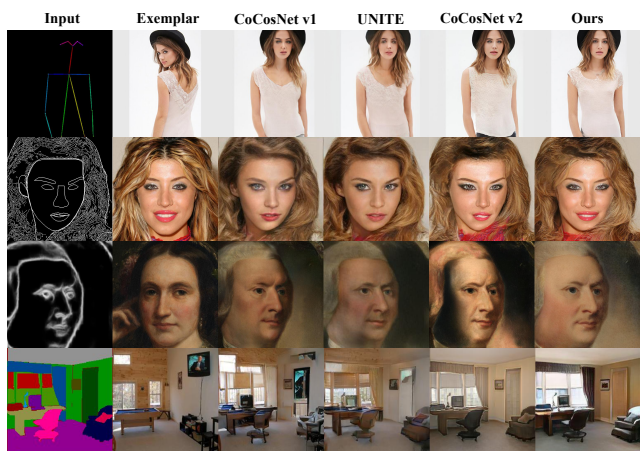


Figure 4: Qualitative comparison results of our method with SOTAs on Deepfashion dataset, CelebA-HQ dataset, MetFaces dataset, and ADE20K dataset, respectively.

have the consistent style and texture as the exemplar. We therefore calculate two additional metrics for evaluating semantic consistency and style similarity. We utilize the pre-trained VGG model on ImageNet and use its high-level feature maps, $relu3_2$, $relu4_2$ and $relu5_2$, to represent high-level semantics. We then calculate the high-level semantics average cosine similarity between the generated and content input images as a semantic consistency metric. Similarly, the style similarity metric can be calculated using the low-level feature maps, $relu1_2$, $relu2_2$.

Implementation Details. We use the TTUR (Heusel et al. 2017) learning strategy with a generator and discriminator learning rate of $1e-4$ and $4e-4$, respectively. We utilize Adam solver with $\beta_1 = 0$ and $\beta_2 = 0.999$. Our experiments are carried out on 8 32GB Tesla V100 GPUs. The input and output feature scales for the CFFT module are $(C, H, W) = (64, 64, 64)$. The weights of the total loss are $\lambda_{align} = 10$, $\lambda_{match} = 10$, $\lambda_{perc} = 0.001$, $\lambda_{cX} = 10$, $\lambda_{adv} = 10$.

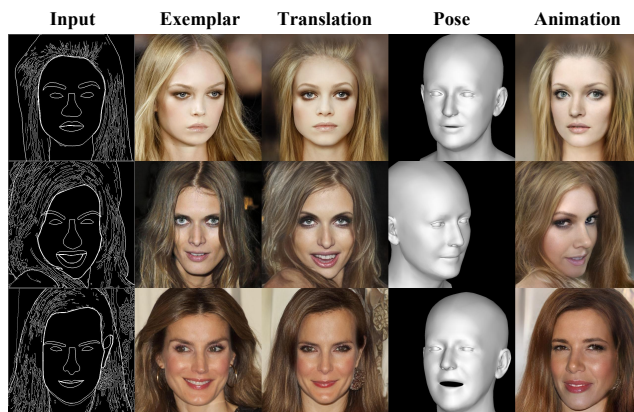


Figure 5: Multi-domain image translation results. Through the figure 3 structure, input edge and exemplar we can get translation result, and then input head pose we can get face animation result based on the translation result.

Experimental Results

Quantitative Evaluation. For the quantitative comparison experiments, we first used FID and SWD to assess the quality of the generated images for each model, and use LPIPS to evaluate the model decoupling and semantic alignment ability on the Deepfashion dataset. As shown in Table 1, our approach outperforms the SOTA methods in most datasets. This is due to the fact that our proposed CFFT module is able to better learn the correspondence between different domains and the potentially beneficial information within each domain, allowing the generator to obtain richer information.

In addition to evaluating the quality of the generated images, we quantitatively assess the semantic consistency of the generated images and content inputs and the style similarity to the exemplar images, as shown in Table 2 and 3. Compared with the SOTA methods, our model also achieves higher performance.

Qualitative Comparison. We make a qualitative comparison with SOTA methods, as shown in Figure 4. Our method




Figure 6: Portrait stylization. Our method can transfer image styles from the MetFaces to the faces of the CelebA-HQ. The model is only trained on the MetFaces dataset. The model input is the edge image of input face.

	FID	SWD	Params
CFFT→CM	14.3	28.8	4.8M
CFFT→AdaIN	16.3	30.2	0.02M
CFFT w/o Hiformer	11.3	25.0	1.37M
CFFT 32×32	9.5	23.7	2.22M
Ours	8.8	19.0	2.62M

Table 4: Ablation studies of our method. Evaluate the FID and SWD metrics on the Deepfashion dataset, and compute the parameters of each module.

can better fuse the style of the exemplar and the semantic content of the input. Our approach can also obtain more refined results when the semantic difference between exemplar and input is large, as shown in the third row. In addition to the comparative results, Figure 1 also show the more diverse generative results of our method. As can be seen from the figures our approach shows strong performance in a wide range of image translation tasks and the diversity of the generated results reflects the generalizability of our model.

Applications. Our method has good potential for application. Firstly, we can decouple and fuse information from multiple domains by cascading CFMT modules. We test a specific task, face animation. As shown in Figure 5, we first learn exemplar-based image translation between edge input and RGB exemplar, and then add head pose to learn the correspondence between three domains. The head pose is extracted via DECA (Feng et al. 2021). The whole training process does not use paired data, and it is self-supervised training. Our generated results have the shape, style, and pose corresponding to the each input domain. By cascading CFMT modules, our approach can progressively learn the feature fusion from multiple image domains to achieve more

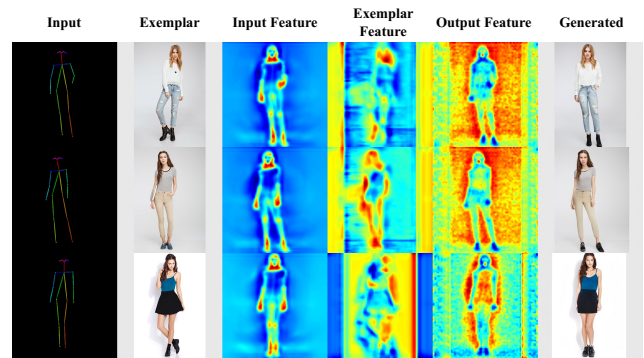


Figure 7: Visualization results of CFMT’s input and output feature maps, which demonstrate the decoupling and semantic alignment capabilities of CFMT.

complex image translation. Secondly, we test the portrait stylization experiment, as shown in Figure 6. We demonstrate the generalizability and excellent performance of our method through this cross-dataset experiment. More comparative results and visualizations are presented in the supplementary material.

Ablation Study

We perform the ablation studies of CFMT, the key module of our model. The quantitative results are shown in Table 4. The first row of the table represents that we replace the CFMT module with the Correlation Matrix in CoCosNet. This method requires a larger amount of parameters and does not achieve good performance. In the second row of the table we replace the CFMT with the AdaIN (Huang and Belongie 2017) structure, which is a very common method of feature transfer. Although it has few parameters, its performance is low for complex correspondence learning. In the third row, we remove the Hiformer in the CFMT, and the performance of the model has dropped significantly. In the fourth row, we decrease the spatial scale of the latent features in CFMT, from 64×64 to 32×32 . But the model is still able to maintain good performance. In addition, we visualise the input and output features of the CFMT module, as shown in Figure 7. The results show that the CFMT is capable of good feature decoupling and semantic alignment.

Conclusion

This paper proposes CFMT-GAN, an exemplar-based image translation that learns the feature fusion of different domains by transformer-based module CFMT, which is able to not only learn the correspondence between domains better but also extract potential beneficial information within each domain. Our proposed CFMT has strong generality, and we can achieve multi-domain image translation by cascading CFMT. Quantitative and qualitative experiments demonstrate the superiority of our approach over the SOTA methods, and our ablation studies on the key module CFMT also reflect its core effect. There is still room for improvement in our approach, e.g., our model still has a performance deficit in the task of mask-to-image. We will optimize this task in future work.

Acknowledgments

This work is supported by the National Key Research and Development Program of China under Grant No. 2021YFC3320103.

References

- Arnab, A.; Deghani, M.; Heigold, G.; Sun, C.; Lučić, M.; and Schmid, C. 2021. Vivit: A video vision transformer. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6836–6846.
- Barnes, C.; Shechtman, E.; Finkelstein, A.; and Goldman, D. B. 2009. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3): 24.
- Cao, Z.; Simon, T.; Wei, S.-E.; and Sheikh, Y. 2017. Realtime multi-person 2d pose estimation using part affinity fields. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7291–7299.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. *European conference on computer vision*, 213–229.
- Choi, Y.; Choi, M.; Kim, M.; Ha, J.-W.; Kim, S.; and Choo, J. 2018. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8789–8797.
- Choi, Y.; Uh, Y.; Yoo, J.; and Ha, J.-W. 2020. Stargan v2: Diverse image synthesis for multiple domains. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8188–8197.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Deghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houslyby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *ICLR*.
- Feng, Y.; Feng, H.; Black, M. J.; and Bolkart, T. 2021. Learning an animatable detailed 3D face model from in-the-wild images. *ACM Transactions on Graphics (TOG)*, 40(4): 1–13.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Huang, X.; and Belongie, S. 2017. Arbitrary style transfer in real-time with adaptive instance normalization. *Proceedings of the IEEE international conference on computer vision*, 1501–1510.
- Huang, X.; Liu, M.-Y.; Belongie, S.; and Kautz, J. 2018. Multimodal unsupervised image-to-image translation. *Proceedings of the European conference on computer vision (ECCV)*, 172–189.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1125–1134.
- Jiang, Y.; Chang, S.; and Wang, Z. 2021. Transgan: Two pure transformers can make one strong gan, and that can scale up. *Advances in Neural Information Processing Systems*, 34.
- Johnson, J.; Alahi, A.; and Fei-Fei, L. 2016. Perceptual losses for real-time style transfer and super-resolution. *European conference on computer vision*, 694–711.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2017. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.
- Karras, T.; Aittala, M.; Hellsten, J.; Laine, S.; Lehtinen, J.; and Aila, T. 2020. Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems*, 33: 12104–12114.
- Lee, H.-Y.; Tseng, H.-Y.; Huang, J.-B.; Singh, M.; and Yang, M.-H. 2018. Diverse image-to-image translation via disentangled representations. *Proceedings of the European conference on computer vision (ECCV)*, 35–51.
- Lee, K.; Chang, H.; Jiang, L.; Zhang, H.; Tu, Z.; and Liu, C. 2021. Vitgan: Training gans with vision transformers. *arXiv preprint arXiv:2107.04589*.
- Li, R.; Jiao, Q.; Cao, W.; Wong, H.-S.; and Wu, S. 2020. Model adaptation: Unsupervised domain adaptation without source data. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9641–9650.
- Li, Y.; Huang, C.; and Loy, C. C. 2019. Dense intrinsic appearance flow for human pose transfer. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3693–3702.
- Liu, Y.; Zhang, W.; and Wang, J. 2021. Source-free domain adaptation for semantic segmentation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1215–1224.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10012–10022.
- Liu, Z.; Luo, P.; Qiu, S.; Wang, X.; and Tang, X. 2016. Deep-fashion: Powering robust clothes recognition and retrieval with rich annotations. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1096–1104.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep learning face attributes in the wild. *Proceedings of the IEEE international conference on computer vision*, 3730–3738.
- Ma, L.; Jia, X.; Georgoulis, S.; Tuytelaars, T.; and Van Gool, L. 2018. Exemplar guided unsupervised image-to-image translation with semantic consistency. *arXiv preprint arXiv:1805.11145*.
- Ma, L.; Jia, X.; Sun, Q.; Schiele, B.; Tuytelaars, T.; and Van Gool, L. 2017. Pose guided person image generation. *Advances in neural information processing systems*, 30.
- Ma, T.; Peng, B.; Wang, W.; and Dong, J. 2021. MUST-GAN: Multi-level Statistics Transfer for Self-driven Person Image Generation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13622–13631.

- Mechrez, R.; Talmi, I.; and Zelnik-Manor, L. 2018. The contextual loss for image transformation with non-aligned data. *Proceedings of the European conference on computer vision (ECCV)*, 768–783.
- Men, Y.; Mao, Y.; Jiang, Y.; Ma, W.-Y.; and Lian, Z. 2020. Controllable person image synthesis with attribute-decomposed gan. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5084–5093.
- Park, T.; Efros, A. A.; Zhang, R.; and Zhu, J.-Y. 2020. Contrastive learning for unpaired image-to-image translation. *European Conference on Computer Vision*, 319–345.
- Park, T.; Liu, M.-Y.; Wang, T.-C.; and Zhu, J.-Y. 2019. Semantic image synthesis with spatially-adaptive normalization. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2337–2346.
- Ren, Y.; Yu, X.; Chen, J.; Li, T. H.; and Li, G. 2020. Deep image spatial transformation for person image generation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7690–7699.
- Saito, K.; Saenko, K.; and Liu, M.-Y. 2020. Coco-funit: Few-shot unsupervised image translation with a content conditioned style encoder. *European Conference on Computer Vision*, 382–398.
- Shen, Y.; Gu, J.; Tang, X.; and Zhou, B. 2020. Interpreting the latent space of gans for semantic face editing. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9243–9252.
- Siarohin, A.; Sangineto, E.; Lathuiliere, S.; and Sebe, N. 2018. Deformable gans for pose-based human image generation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3408–3416.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, H.; Zhu, Y.; Green, B.; Adam, H.; Yuille, A.; and Chen, L.-C. 2020. Axial-deeplab: Stand-alone axial-attention for panoptic segmentation. *European Conference on Computer Vision*, 108–126.
- Wang, M.; Yang, G.-Y.; Li, R.; Liang, R.-Z.; Zhang, S.-H.; Hall, P. M.; and Hu, S.-M. 2019. Example-guided style-consistent image synthesis from semantic labeling. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1495–1504.
- Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018. High-resolution image synthesis and semantic manipulation with conditional gans. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8798–8807.
- Wu, R.; Zhang, G.; Lu, S.; and Chen, T. 2020. Cascade efgan: Progressive facial expression editing with local focuses. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5021–5030.
- Xie, S.; and Tu, Z. 2015. Holistically-nested edge detection. *Proceedings of the IEEE international conference on computer vision*, 1395–1403.
- Yi, Z.; Zhang, H.; Tan, P.; and Gong, M. 2017. Dualgan: Unsupervised dual learning for image-to-image translation. *Proceedings of the IEEE international conference on computer vision*, 2849–2857.
- Yu, J.; Lin, Z.; Yang, J.; Shen, X.; Lu, X.; and Huang, T. S. 2018. Generative image inpainting with contextual attention. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5505–5514.
- Zhan, F.; Yu, Y.; Cui, K.; Zhang, G.; Lu, S.; Pan, J.; Zhang, C.; Ma, F.; Xie, X.; and Miao, C. 2021. Unbalanced feature transport for exemplar-based image translation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15028–15038.
- Zhang, P.; Zhang, B.; Chen, D.; Yuan, L.; and Wen, F. 2020. Cross-domain correspondence learning for exemplar-based image translation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5143–5153.
- Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene parsing through ade20k dataset. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 633–641.
- Zhou, X.; Zhang, B.; Zhang, T.; Zhang, P.; Bao, J.; Chen, D.; Zhang, Z.; and Wen, F. 2021. Cocosnet v2: Full-resolution correspondence learning for image translation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11465–11475.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017a. Unpaired image-to-image translation using cycle-consistent adversarial networks. *Proceedings of the IEEE international conference on computer vision*, 2223–2232.
- Zhu, J.-Y.; Zhang, R.; Pathak, D.; Darrell, T.; Efros, A. A.; Wang, O.; and Shechtman, E. 2017b. Toward multimodal image-to-image translation. *Advances in neural information processing systems*, 30.