

DOMAIN ADAPTATION FOR LEARNING GENERATOR FROM PAIRED FEW-SHOT DATA

Chun-Chih Teng[†] Pin-Yu Chen[‡] Wei-Chen Chiu[†]

[†]National Chiao Tung University [‡]IBM Research

ABSTRACT

We propose a Paired Few-shot GAN (PFS-GAN) model for learning generators with sufficient source data and a few target data. While generative model learning typically needs large-scale training data, our PFS-GAN not only uses the concept of few-shot learning but also domain shift to transfer the knowledge across domains, which alleviates the issue of obtaining low-quality generator when only trained with target domain data. The cross-domain datasets are assumed to have two properties: (1) each target-domain sample has its source-domain correspondence and (2) two domains share similar content information but different appearance. Our PFS-GAN aims to learn the disentangled representation from images, which composed of domain-invariant content features and domain-specific appearance features. Furthermore, a relation loss is introduced on the content features while shifting the appearance features to increase the structural diversity. Extensive experiments show that our method has better quantitative and qualitative results on the generated target-domain data with higher diversity in comparison to several baselines.

Index Terms— Few-Shot Learning, Domain Adaptation, Generative Adversarial Networks

1. INTRODUCTION

Generative Adversarial Network (GAN) [1] is one of the most popular generative models based on deep learning nowadays, in which its learning process typically relies on sufficient amount of training data in order to generate realistic and diverse data samples. However, in some tasks such as medical applications, the training data are often scarce and the cost of acquiring new data samples is expensive, if not impossible. Consequently, the restriction of limited training data size may largely degrade the performance of GAN thus leading to mode collapse. To overcome this challenge, we propose to use domain adaptation [2, 3, 4, 5] to transfer the knowledge of generative process from another data collection with rich information (*source domain*) to the target task (*target domain*), where these two domains share some similarities but follow two different data distributions, also known as *domain shift*. While domain adaptation has mostly been used for




Fig. 1: Illustration of the paired few-shot (PFS) learning and generation problem setup in this paper. There are abundant and diverse training data in the source domain but only few training data in the target domain with known correspondence to their source-domain counterparts. Our PFS-GAN uses these cross-domain data and paired few-shot information to train a target-domain generator.

classification [6, 7, 8], its application on learning generators with limited data is still challenging and not well-explored. It is also worth noting that: Image-to-image translation (e.g., CycleGAN [9]) enables cross-domain data translation but does not address data generation; and other GAN models (e.g., CoGAN [10]) which can generate multi-domain images simultaneously still require a large amount of data to learn the satisfactory target-domain generator.

As shown in Figure 1, this paper aims to learn a target-domain generator with sufficient source-domain data but only few paired target-domain data. Specifically, we introduce few-shot learning to help transferring the knowledge of generation from source to target domain. Noting that since most existing few-shot learning works focus on classification and use supervised metrics for learning new classes, they are not directly applicable to GAN. To integrate the concept of few-shot learning into learning target-domain generator, we consider the *paired few-shot setting* where each target-domain sample has a correspondence in the source, which means these two domains share similar contents but have difference appearances. Under this setting, our proposed **Paired Few-Shot GAN (PFS-GAN)** could transfer the knowledge across domains by domain adaptation, while learning to generate target-domain samples with high diversity and fidelity. Particularly, the cross-domain knowledge is learned via disentangled representations composed of domain-invariant content features and domain-specific appearance features. While the target-domain generator learns the appearance feature from the few target-domain data, the rich content information transferred from the source domain can thus help improving the diversity of structure generation in the target domain.

National Chiao Tung University and National Yang-Ming University are merged into National Yang Ming Chiao Tung University in February 2021.

2. PROPOSED METHOD: PFS-GAN

Given a cross-domain dataset built upon the source-domain $X_{src} = \{x_{src,i}\}_{i=1}^{N_{src}}$ and target-domain $X_{tar} = \{x_{tar,i}\}_{i=1}^{N_{tar}}$ samples where $N_{tar} \ll N_{src}$, our problem scenario assumes that for each $x_{tar,i}$ there is a paired source-domain sample $x_{src,\kappa(i)}$ in which $x_{tar,i}$ and $x_{src,\kappa(i)}$ have the similar content/structure but different appearance/texture, where $\kappa(\cdot)$ is a one-to-one mapping function to obtain the index correspondence across domains. All these cross-domain pairs are denoted as $X^{Pair} = \{x_{tar,i}, x_{src,\kappa(i)}\}_{i=1}^{N_{tar}}$. In order to transfer the richer content from the source to the target domain and make the target-domain generator capable of generating target-domain samples with more diverse content, we first factor out the content information from the source-domain samples (i.e. **disentanglement**), then tackle the challenges of domain adaptation and few-shot setting to benefit the target-domain generator. Our proposed PFS-GAN consists of generators $\{G_{src}, G_{tar}\}$, discriminators $\{D_{src}, D_{tar}\}$, source content encoder E_{src}^C , appearance encoders $\{E_{src}^A, E_{tar}^A\}$ and the relation network R . We now detail the two training stages for PFS-GAN (as depicted in Figure 2) in the following.

2.1. First stage: learning G_{src} and disentanglement

The disentanglement of the source-domain latent space into content and appearance parts is realized by a framework based on VAE [11] and GAN. Given a source-domain sample x_{src} , source-domain content encoder E_{src}^C and appearance encoder E_{src}^A map it into the content feature z_{src}^C and the appearance feature z_{src}^A respectively, where $\{z_{src}^C, z_{src}^A\}$ can be used to reconstruct x_{src} by source-domain generator G_{src} with their distribution modeled by standard normal distribution $\mathcal{N}(0, I)$. The source-domain image reconstruction loss \mathcal{L}_{src}^{IR} and the KL-divergence loss \mathcal{L}_{src}^{KL} are thus defined as:

$$\begin{aligned}\mathcal{L}_{src}^{IR} &= \mathbb{E}_{x_{src} \sim X_{src}} \|G_{src}(z_{src}^C, z_{src}^A) - x_{src}\| \\ \mathcal{L}_{src}^{KL} &= \mathbb{E}_{x_{src} \sim X_{src}} [D_{KL}(E_{src}^C(x_{src}) || \mathcal{N}(0, I))] \\ &\quad + \mathbb{E}_{x_{src} \sim X_{src}} [D_{KL}(E_{src}^A(x_{src}) || \mathcal{N}(0, I))]\end{aligned}$$

For better encouraging the disentanglement between the content and appearance features, two objectives are introduced additionally. First, given two images generated by the same content vector z_{src}^C but different appearance vectors (z_{src}^A and z_{src}^A'), they should have high similarity on the content/structure. Since the higher layers of the ImageNet-pretrained VGG network tend to represent the high-level structure content of an image [12], we use the perceptual loss \mathcal{L}^P [13] to ensure z_{src}^C maintaining the content information:

$$\mathcal{L}^P = \left\| F(G(z_{src}^C, z_{src}^A)) - F(G(z_{src}^C, z_{src}^A')) \right\|$$

where F is pretrained VGG network up to relu3.3 layer. Second, for preventing appearance vector z_{src}^A from being ignored by the generator G_{src} , we have the latent reconstruction loss \mathcal{L}_{src}^{AR} on the appearance feature z_{src}^A :

$$\mathcal{L}_{src}^{AR} = \mathbb{E}_{z_{src}^A \sim \mathcal{N}(0, I)} \|E_{src}^A(G_{src}(z_{src}^C, z_{src}^A)) - z_{src}^A\|$$

in which it means that for a synthetic image $G_{src}(z_{src}^C, z_{src}^A)$, we should be able to obtain z_{src}^A from it via the appearance encoder E_{src}^A . Note that here $z_{src}^A \sim \mathcal{N}(0, I)$ and z_{src}^C can be any content vector extracted from the source-domain samples. Moreover, we improve the realness of the synthetic \tilde{X}_{src} via adversarial learning (based on the hinge loss [14]), where

$$\begin{aligned}\mathcal{L}_{src}^{IA, D} &= \mathbb{E}_{x_{src} \sim X_{src}} [1 - D_{src}(x_{src})]_+ \\ &\quad + \mathbb{E}_{\tilde{x}_{src} \sim \tilde{X}_{src}} [1 + D_{src}(\tilde{x}_{src})]_+ \\ \mathcal{L}_{src}^{IA, G} &= - \mathbb{E}_{\tilde{x}_{src} \sim \tilde{X}_{src}} [D_{src}(\tilde{x}_{src})]\end{aligned}$$

are used to update D_{src} (discriminator) and G_{src} respectively.

2.2. Second stage: learning target generator G_{tar}

Once the source-domain generator G_{src} and its latent space disentanglement are learned, we now aim to train the target-domain generator G_{tar} which combines the diverse content inherited from the source with the target-domain specific appearance to synthesize target-domain samples. To this end, we also require the disentanglement of target-domain latent space but now it is realized by $\{E_{src}^C, E_{tar}^A\}$ and G_{tar} (noting that we use E_{src}^C here as two domains share the same content). First, as a target-domain sample $x_{tar,i}$ shares the same content information with its source-domain correspondence $x_{src,\kappa(i)}$, the image reconstruction loss \mathcal{L}_{tar}^{IR} is defined as:

$$\mathcal{L}_{tar}^{IR} = \sum_{i=1}^{N_{tar}} \|G_{tar}(E_{src}^C(x_{src,\kappa(i)}), E_{tar}^A(x_{tar,i})) - x_{tar,i}\|$$

Second, we would also like to regularize the distribution of target-domain appearance features z_{tar}^A by $\mathcal{N}(0, I)$. However, as there are only few-shots for target-domain data which could be problematic to directly apply the regularization on their distribution, we adopt the data augmentation to increase the appearance variation of target-domain data, where the augmentation is performed by randomly shifting the chromatic channels in the Lab color space. The KL-divergence loss \mathcal{L}_{tar}^{KL} is defined on the augmented data samples X_{tar}^{aug} :

$$\mathcal{L}_{tar}^{KL} = \mathbb{E}_{x_{tar} \sim X_{tar}^{aug}} [D_{KL}(E_{tar}^A(x_{tar}) || \mathcal{N}(0, I))]$$

Third, we also use the latent reconstruction loss \mathcal{L}_{tar}^{AR} to encourage the enhance the disentanglement:

$$\mathcal{L}_{tar}^{AR} = \mathbb{E}_{z_{tar}^A \sim E_{tar}^A(X_{tar}^{aug})} \|E_{tar}^A(G_{tar}(z_{src}^C, z_{tar}^A)) - z_{tar}^A\|$$

where now z_{tar}^A is sampled from $E_{tar}^A(X_{tar}^{aug})$ and z_{src}^C can be any content vector extracted from $x_{src} \in X^{Pair}$. Fourth, we also impose adversarial loss on the synthetic target-domain samples X_{tar}^{syn} which use source content $z_{src}^C \sim E_{src}^C(X_{src})$ and target appearance features $z_{tar}^A \sim E_{tar}^A(X_{tar})$, where

$$\begin{aligned}\mathcal{L}_{tar}^{IA, D} &= \mathbb{E}_{x_{tar} \sim X_{tar}} [1 - D_{tar}(x_{tar})]_+ \\ &\quad + \mathbb{E}_{x_{tar} \sim X_{tar}^{syn}} [1 + D_{tar}(x_{tar})]_+ \\ \mathcal{L}_{tar}^{IA, G} &= - \mathbb{E}_{x_{tar} \sim X_{tar}^{syn}} [D_{tar}(x_{tar})]\end{aligned}$$

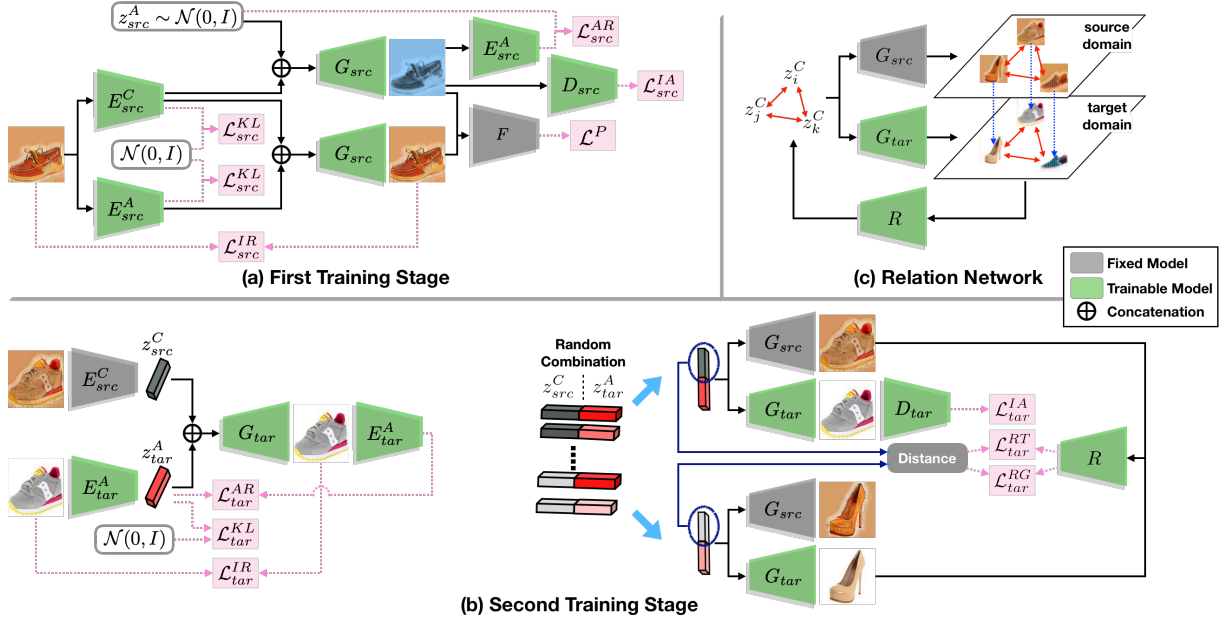


Fig. 2: Overview of our proposed Paired Few-Shot GAN (PFS-GAN) model. (a) The first training stage: learning source-domain generator and disentanglement (cf. Section 2.1). (b) The second training stage: learning target-domain generator (cf. Section 2.2). (c) Relation network leverages the content relation between cross-domain samples for regularizing the content diversity between generated target-domain samples.

are used to update D_{tar} (discriminator) and G_{tar} respectively.

Lastly, in order to better enforce the generated target-domain samples on having the rich content adapted from the source, we particularly propose a novel **Relation Loss** which leverages the relation among content vectors of cross-domain samples to regularize the content diversity between generated target-domain samples, where the idea is shown in Figure 2(c): the relations between different content vectors z^C (i.e. pairwise distance) should be well reflected on the content similarity among the related generated samples, which we would like to have such property in both source and target domains. To this end, we first define the content similarity between cross-domain images: Given any cross-domain image pair $\{x_{src,j}, x_{tar,i}\} \notin X^{Pair}$ where $j \neq \kappa(i)$, their content similarity is computed by

$$D_c(x_{src,j}, x_{tar,i}) = \|E_{src}^C(x_{src,j}) - E_{src}^C(x_{src,\kappa(i)})\|$$

as $x_{tar,i}$ and its paired $x_{src,\kappa(i)}$ should have the same content information. We then train a relation network R learning to project the cross-domain image pair into a value representing their content similarity, which exactly equals to the L2 distance between their corresponding content vectors, with the objective:

$$\mathcal{L}_{tar}^{RT} = \sum_i^{N_{tar}} \sum_{j \neq \kappa(i)}^{N_{src}} \|R(x_{src,j}, x_{tar,i}) - D_c(x_{src,j}, x_{tar,i})\|$$

Once the relation network R is learnt, we use it to regularize the content diversity between the generated target-domain samples: Given a cross-domain pair of generated samples, $\{\tilde{x}_{src} = G_{src}(z_{src,i}^C, z_{src,i}^A), \tilde{x}_{tar} = G_{tar}(z_{src,j}^C, z_{tar,i}^A)\}$ where

$i \neq j$, $z_{src,i}^C$ and $z_{src,j}^C$ are taken from the source-domain images by E_{src}^C , $z_{src,i}^A \sim \mathcal{N}(0, I)$, and $z_{tar,i}^A \in E_{tar}^A(X_{tar})$, their content similarity computed by R should be equal to the L2 distance between $z_{src,i}^C$ and $z_{src,j}^C$, leading to the relation loss:

$$\mathcal{L}_{tar}^{RG} = \mathbb{E} \|R(\tilde{x}_{src}, \tilde{x}_{tar}) - D_c(\tilde{x}_{src}, \tilde{x}_{tar})\|$$

where we sample many cross-domain pairs of $\{\tilde{x}_{src}, \tilde{x}_{tar}\}$. The gradient of \mathcal{L}_{tar}^{RG} helps update G_{tar} and enforce it to generate target-domain samples with sufficient content diversity.

3. EXPERIMENTS

























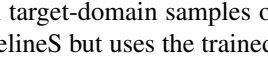

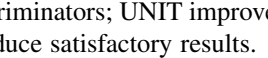
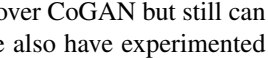
Datasets Three datasets, **CUFS** [17], **UT Zappos50K** [15], and **Facescrub** [16], are used to build four cross-domain experimental settings. UT Zappos50K is a large dataset of shoes images with sketches provided for each real shoes image; Facescrub dataset is composed of face images of 530 celebrities, taken under various poses/angles; CUFS dataset consists of image pairs of face portraits and their corresponding sketches. Four experimental settings are listed as follows: (1) **Edge2Shoes** picks 10 real shoes images as the target-domain, while their corresponding sketches with many others (in total 10,000 sketches) are taken as the source-domain; (2) **Style2Shoes** is almost identical to Edge2Shoes but the sketches are stylized by [18]; (3) **Face2Face** is based on Facescrub, where the source-domain is composed of photos of 525 celebrities while the target-domain has 10 photos from each of the left 5 celebrities (i.e. adapting diverse poses from source to target); (4) **Sketch2Face** randomly selects 10 real portraits as target-domain while having all 1,194 sketches as source.

Baselines Four baselines are used for comparison with our proposed PFS-GAN: (1) **BaselineS** is a GAN trained from

Table 1: Quantitative comparison of target-domain generation among various approaches under different experimental settings and metrics.

Experiment Synthesis Manner	Edge2Shoes (E2S)				Style2Shoes (S2S)				Face2Face (F2F)				Sketch2Face (S2F)			
	Rand		Syn		Rand		Syn		Rand		Syn		Rand		Syn	
Metric	FID	KID	FID	KID	FID	KID	FID	KID	FID	KID	FID	KID	FID	KID	FID	KID
BaselineS	185.6	0.37	-	-	219.9	0.18	-	-	143.8	0.19	-	-	122.6	0.20	-	-
BaselineT	165.0	0.29	-	-	149.9	0.14	-	-	111.4	0.22	-	-	125.0	0.19	-	-
CoGAN	197.2	0.38	-	-	179.0	0.14	-	-	220.2	0.15	-	-	195.7	0.35	-	-
UNIT	135.8	0.34	-	-	143.3	0.27	-	-	103.1	0.42	-	-	93.04	0.44	-	-
PFS-GAN [‡]	125.9	0.21	103.5	0.12	116.1	0.13	93.37	0.17	55.47	0.08	32.67	0.04	48.93	0.08	44.83	0.06
PFS-GAN [†]	91.42	0.19	88.93	0.08	96.08	0.11	78.83	0.12	50.17	0.06	29.45	0.03	45.61	0.06	40.78	0.06
PFS-GAN	92.55	0.17	83.31	0.07	79.94	0.08	67.22	0.11	40.76	0.04	25.11	0.02	38.50	0.04	36.95	0.04

Table 2: Qualitative examples produced by various baselines and our proposed PFS-GAN model under different experimental settings.

	Style2Shoes (S2S)	Edge2Shoes (E2S)	Face2Face (F2F)	Sketch2Face (S2F)
Source Data				
Target Data				
BaselineS				
BaselineT				
CoGAN				
UNIT				
PFS-GAN				

scratch based on target-domain samples only; (2) **BaselineT** is similar to BaselineS but uses the trained G_{src} as its initialization; (3) **CoGAN** [9] consists of two tuple of GAN trained by the cross-domain data with the weights of their networks partially shared; (4) **UNIT** [3] is the extension of CoGAN which is also trained by the cross-domain data, with having additional loss to regularize the paired cross-domain images being mapped into the same latent vector by the encoders.

Fréchet Inception distance (FID [20]) and Kernel Inception Distance (KID [21]) are used (both the lower the better) for quantitative evaluation on the target-domain generators. Two PFS-GAN variants are also used to perform the ablation study: (1) **PFS-GAN[†]** removes the relation loss and (2) **PFS-GAN[‡]** further removes the adversarial loss in the second training stage. Two manners are used for PFS-GAN to generate target-domain samples: **Rand** takes both z^C and z^A sampled from $\mathcal{N}(0, I)$ while **Syn** takes $\{z^C, z^A\}$ sampled from $\{E_{src}^C(X_{src}), E_{tar}^A(X_{tar})\}$ respectively. Quantitative results are shown in Table 1, where our full PFS-GAN clearly has superior performance than the baselines, and the comparison to the variants verifies the contributions of our designed objectives (particularly the relation loss). Moreover, the difference between PFS-GAN[‡] and the baselines shows the importance of disentanglement. BaselineS and BaselineT perform around the second worst, since they are trained only with few-shot target-domain samples hence easily leading to mode collapse; CoGAN mostly performs the worst as the number of samples for both domains is quite unbalanced thus leading to unsta-

ble discriminators; UNIT improves over CoGAN but still can not produce satisfactory results. We also have experimented on increasing target-domain samples from 10 to 20 then 100. Take Sketch2Face as an example, PFS-GAN boosts from 38.5 to 33.6 then achieves 28.4 in FID, while other baselines are also improved but still worse than ours by a margin with the issues of unstable training and mode collapse remained.

Table 2 provides the qualitative examples. We observe that: BaselineS, BaselineT, and CoGAN generate realistic images, but clearly attempts to memorize the training samples and lose the content diversity (i.e. mode collapse caused by few-shot data). UNIT also suffers from mode collapse (e.g. Style2Shoes and Face2Face), while occasionally produces results with poor quality as in the Sketch2Face case. Our PFS-GAN provides favorable results on both content diversity and fidelity, and is able to synthesize images with having the content that is never seen in the target-domain training samples.

4. CONCLUSION

We propose PFS-GAN to tackle the generative model learning with cross-domain data, where the target-domain has only few-shots provided with paired source-domain samples. PFS-GAN combines the learning disentanglement (domain-invariant content and domain-specific appearance features), domain adaptation, and the cross-domain relation built upon the properties of training data. The target-domain generator from PFS-GAN experimentally shows its capacity on improving the content diversity of generated images and providing superior performance in comparison to several baselines.

5. REFERENCES

- [1] 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 (NeurIPS)*, 2014.
- [2] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al., “Matching networks for one shot learning,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [3] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [4] Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip HS Torr, and Timothy M Hospedales, “Learning to compare: Relation network for few-shot learning,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [5] Yongqin Xian, Saurabh Sharma, Bernt Schiele, and Zeynep Akata, “f-vaegan-d2: A feature generating framework for any-shot learning,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10275–10284.
- [6] Saeid Motiian, Quinn Jones, Seyed Iranmanesh, and Gianfranco Doretto, “Few-shot adversarial domain adaptation,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [7] Feng Liu, Jie Lu, Bo Han, Gang Niu, Guangquan Zhang, and Masashi Sugiyama, “Butterfly: Robust one-step approach towards wildly-unsupervised domain adaptation,” *ArXiv:1905.07720*, 2019.
- [8] Shuhan Tan, Jiening Jiao, and Wei-Shi Zheng, “Weakly supervised open-set domain adaptation by dual-domain collaboration,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [9] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros, “Image-to-image translation with conditional adversarial networks,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [10] Ming-Yu Liu and Oncel Tuzel, “Coupled generative adversarial networks,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [11] Diederik P Kingma and Max Welling, “Auto-encoding variational bayes,” in *International Conference on Learning Representations (ICLR)*, 2014.
- [12] Leon A Gatys, Alexander S Ecker, and Matthias Bethge, “Image style transfer using convolutional neural networks,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [13] Justin Johnson, Alexandre Alahi, and Li Fei-Fei, “Perceptual losses for real-time style transfer and super-resolution,” in *European Conference on Computer Vision (ECCV)*, 2016.
- [14] Dustin Tran, Rajesh Ranganath, and David Blei, “Hierarchical implicit models and likelihood-free variational inference,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [15] A. Yu and K. Grauman, “Semantic jitter: Dense supervision for visual comparisons via synthetic images,” in *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [16] Hong-Wei Ng and Stefan Winkler, “A data-driven approach to cleaning large face datasets,” in *IEEE International Conference on Image Processing (ICIP)*, 2014.
- [17] Xiaogang Wang and Xiaoou Tang, “Face photo-sketch synthesis and recognition,” *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2008.
- [18] Xun Huang and Serge Belongie, “Arbitrary style transfer in real-time with adaptive instance normalization,” in *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [19] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther, “Autoencoding beyond pixels using a learned similarity metric,” in *International Conference on Machine Learning (ICML)*, 2016.
- [20] 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 (NeurIPS)*, 2017.
- [21] Mikołaj Bińkowski, Dougal J Sutherland, Michael Arbel, and Arthur Gretton, “Demystifying mmd gans,” in *International Conference on Learning Representations (ICLR)*, 2018.