

# DynaGAN: Dynamic Few-shot Adaptation of GANs to Multiple Domains

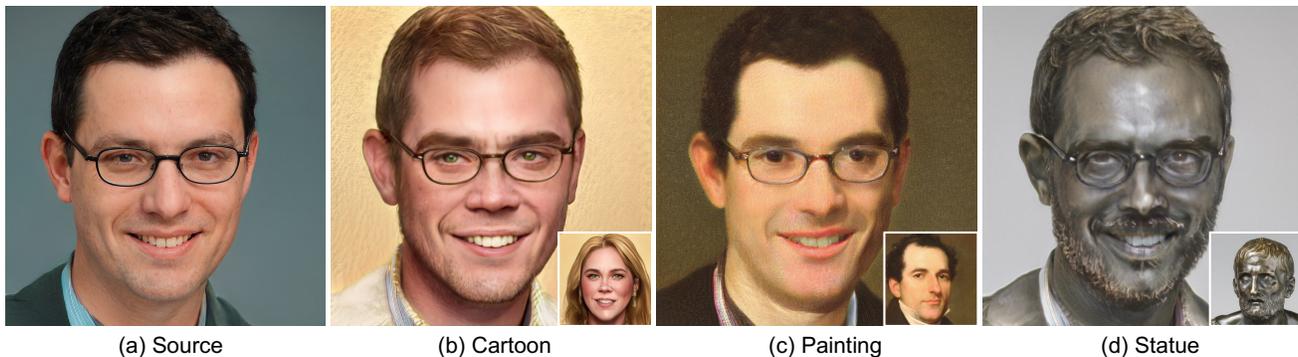
Seongtae Kim  
POSTECH  
South Korea  
seongtae0205@postech.ac.kr

Kyoungkook Kang  
POSTECH  
South Korea  
kkang831@postech.ac.kr

Geonung Kim  
POSTECH  
South Korea  
k2woong92@postech.ac.kr

Seung-Hwan Baek  
POSTECH  
South Korea  
shwbaek@postech.ac.kr

Sunghyun Cho  
POSTECH  
South Korea  
Pebblous  
South Korea  
s.cho@postech.ac.kr



**Figure 1:** We present DynaGAN, a method that dynamically adapts a generator to each target domain for synthesizing multi-domain images using a single generator. (a) For a source image and ten target images in different domains (three of them shown in insets), we demonstrate successful domain-adaptation results: (b) cartoon, (c) painting, and (d) statue. Inputs in (c) and (d): The Metropolitan Museum of Art [Public Domain].

## ABSTRACT

Few-shot domain adaptation to multiple domains aims to learn a complex image distribution across multiple domains from a few training images. A naïve solution here is to train a separate model for each domain using few-shot domain adaptation methods. Unfortunately, this approach mandates linearly-scaled computational resources both in memory and computation time and, more importantly, such separate models cannot exploit the shared knowledge between target domains. In this paper, we propose DynaGAN, a novel few-shot domain-adaptation method for multiple target domains. DynaGAN has an adaptation module, which is a hyper-network that dynamically adapts a pretrained GAN model into the multiple target domains. Hence, we can fully exploit the shared

knowledge across target domains and avoid the linearly-scaled computational requirements. As it is still computationally challenging to adapt a large-size GAN model, we design our adaptation module to be lightweight using the rank-1 tensor decomposition. Lastly, we propose a contrastive-adaptation loss suitable for multi-domain few-shot adaptation. We validate the effectiveness of our method through extensive qualitative and quantitative evaluations.

## CCS CONCEPTS

• **Computing methodologies** → *Image processing; Artificial intelligence.*

## KEYWORDS

Generative adversarial networks, domain adaptation, few-shot learning, hyper-networks

## ACM Reference Format:

Seongtae Kim, Kyoungkook Kang, Geonung Kim, Seung-Hwan Baek, and Sunghyun Cho. 2022. DynaGAN: Dynamic Few-shot Adaptation of GANs to Multiple Domains. In *SIGGRAPH Asia 2022 Conference Papers (SA '22 Conference Papers)*, December 6–9, 2022, Daegu, Republic of Korea. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3550469.3555416>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*SA '22 Conference Papers, December 6–9, 2022, Daegu, Republic of Korea*

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9470-3/22/12...\$15.00

<https://doi.org/10.1145/3550469.3555416>

## 1 INTRODUCTION

Recent progress of generative adversarial networks (GANs) has opened up a new chapter in image synthesis [Brock et al. 2019; Goodfellow et al. 2014; Karras et al. 2021, 2019, 2020b], extending its application to various tasks including data augmentation [Sandfort et al. 2019], image restoration [Chan et al. 2021; Wang et al. 2021; Yang et al. 2021], and image and video manipulation [Kang et al. 2021; Patashnik et al. 2021; Richardson et al. 2021; Shen et al. 2020; Shen and Zhou 2021; Tov et al. 2021; Tzaban et al. 2022; Zhu et al. 2020]. Unfortunately, this success has been mainly demonstrated on a few large-scale datasets such as human portraits [Karras et al. 2019; Liu et al. 2015], because of the fundamental requirement of GANs on many training samples. Thus, it is still challenging to apply GANs to diverse domains where data is scarce. While recent methods [Jiang et al. 2021; Karras et al. 2020a; Zhao et al. 2020] have alleviated the burden of data-intense GAN training, they still require hundreds of training images.

To circumvent collecting a large-scale dataset, domain adaptation methods have been proposed [Noguchi and Harada 2019; Wang et al. 2020, 2018]. Their core idea is to exploit a GAN model pretrained on a data-sufficient domain and finetune it onto a data-scarce domain where common knowledge exists between the two domains. Recent few-shot domain adaptation methods have demonstrated impressive success with less than ten training images [Gal et al. 2022; Ojha et al. 2021; Zhu et al. 2022]. However, they assume that training images have similar appearances, and fail to handle training images with diverse domains as shown in Fig. 4, generating an *average style* of multiple domains.

One can bypass such failure by learning a separate model for each target domain, resulting in many models depending on the number of target domains. Unfortunately, this approach mandates linearly-scaled computational resources both in memory and computation time. More importantly, such separate models cannot exploit the shared knowledge between target domains, e.g. common structural composition in oil-painting portraits by different artists. This results in the limited capability of domain-adapted models.

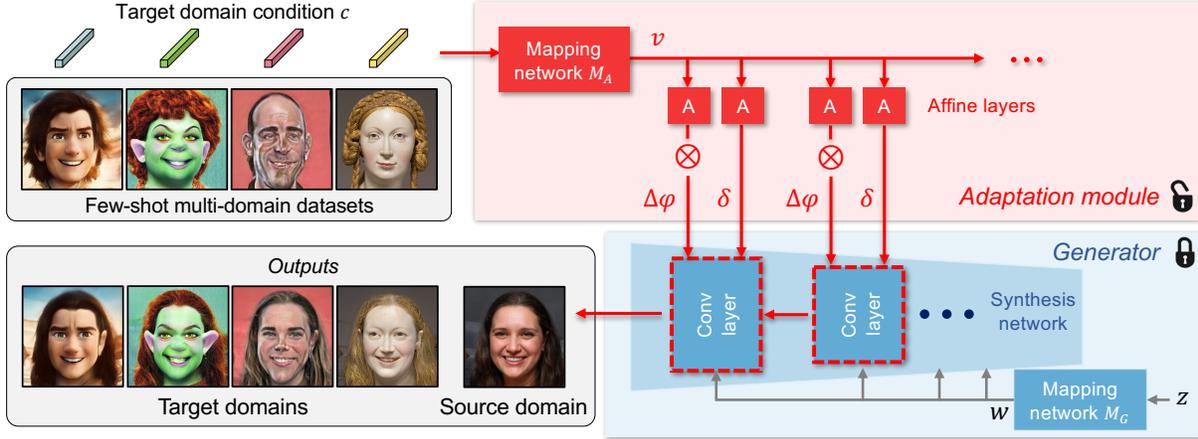
This paper proposes DynaGAN, a novel few-shot domain adaptation method to handle multiple target domains. We depart from the separate-model principle and use a single GAN model pretrained on a data-sufficient domain. To this end, we propose an adaptation module, a *hyper-network* that dynamically adapts a GAN model to each target domain. This change allows us to avoid the problems of computational resources and overlooked shared knowledge between target domains. Specifically, our adaptation module takes a target-domain condition vector as an input and modulates the weights of a GAN model, which acts as a dynamic domain adaptation to a target domain. We carefully design our adaptation module to be light-weight and memory-efficient with rank-1 tensor decomposition. We train the adaptation module with a training dataset consisting of multiple target domains, each of which may have extremely few-shot images, e.g., only one image. As we aim to maintain the unique style of each target domain, we further introduce a contrastive-adaptation loss to preserve the distinctive attributes of different target domains. As shown in Fig. 1, DynaGAN allows us to synthesize images of diverse target domains, e.g., target domains of different drawing styles, and different animal species.

In addition to the reduced computational cost and diverse target domain adaptation, our method enjoys additional benefits. First, our approach performs domain adaptation using an adaptation module that modulates only the weights of the convolutional layers of a GAN model while preserving the GAN model itself intact. As a result, the semantic editability of the original GAN model is naturally maintained in DynaGAN, where the degree of adaptation can be easily controlled by scaling the amount of modulation from the adaptation module. Second, as our adaptation module is trained on multi-domain images at once, it can effectively learn the shared knowledge across different domains, hence avoiding overfitting. Third, our approach allows interpolation of target-domain condition vectors, meaning that we can synthesize an image in a new target domain. Our extensive experiments show that DynaGAN achieves superior performance compared to the existing methods.

## 2 RELATED WORK

*Generative Adversarial Networks.* Since GANs were firstly introduced by Goodfellow *et al.* [2014], synthesizing realistic images has been one of the core tasks and test-beds for GAN models. While we have witnessed the remarkable success of recent GAN-based synthesis methods including StyleGAN variants [Karras et al. 2021, 2019, 2020b], one pitfall is that training a GAN model mandates a large-scale dataset, e.g., human portraits in FFHQ [Karras et al. 2019] and church and car images in LSUN [Yu et al. 2015]. Data augmentation is a promising way for overcoming the large-scale data problem by effectively increasing the training samples based on certain augmentation rules [Jiang et al. 2021; Karras et al. 2020a; Zhao et al. 2020]. Unfortunately, state-of-the-art data-augmentation methods still require at least hundreds of images.

*Domain Adaptation for GANs.* Domain adaptation is another way of detouring the large-scale data requirement of GANs [Huang et al. 2022; Wang et al. 2020, 2018]. Once a GAN model is pretrained on a data-sufficient domain, domain-adaptation methods finetune the model on a data-scarce domain. One of the major challenges of domain adaptation for GANs is that a GAN model easily overfits a few samples in the data-scarce domain. This results in severe mode-collapse artifacts, meaning that an adapted GAN model can only synthesize virtually identical images [Wang et al. 2018]. Thus, few-shot domain adaptation of GANs has remained a challenging problem. Recent methods have tackled this by means of target-weighted GAN latent space [Wang et al. 2020], adapting batch statistics [Noguchi and Harada 2019], and regularizing the change of weights in a GAN model [Li et al. 2020]. Ojha *et al.* [2021] propose to maintain the diversity of domain-adapted images as in the original GAN model by imposing a cross-domain consistency loss. StyleGAN-nada [Gal et al. 2022] exploits the recently-proposed CLIP representation [Radford et al. 2021], which is a learned prior of natural images and text, in order to linearly translate the latent space from a source to a target domain maintaining the source diversity. Mind-the-gap (MTG) [Zhu et al. 2022] takes a step forward to achieve one-shot domain adaptation by preserving the source-domain diversity with reference-based CLIP regularizations. Unfortunately, the existing domain-adaptation methods fail to handle a multi-domain few-shot dataset, for which our method is designed.



**Figure 2: DynaGAN network consists of a generator pretrained on a source domain and an adaptation module that dynamically modulates the generator parameters. DynaGAN takes a one-hot vector  $c$  encoding a target domain as an input, projects it to a continuous representation  $v$  through a mapping network  $M_A$ , and estimates modulation parameters  $\Delta\varphi$  and  $\delta$  through affine layers. These parameters modulate the weight of each convolutional layer in the generator. As a result, the generator is adapted to the target domain. 4th image in the few-shot multi-domain datasets: The Metropolitan Museum of Art [Public Domain].**

*Hyper-networks.* Hyper-networks, firstly proposed by Ha and Le [2017], are neural networks that modulate the parameters of other neural networks. Offering generalization and flexibility to existing neural networks, hyper-networks are being actively explored in many applications including neural architecture search [Ratzlaff and Fuxin 2019; Zoph et al. 2018], multi-task learning [Mahabadi et al. 2021], continual learning [von Oswald et al. 2020], semantic segmentation [Nirkin et al. 2021], and 3D modeling [Littwin and Wolf 2019]. Most relevant to us, Alaluf *et al.* [2022] and Dinh *et al.* [2022] demonstrate that using a hyper-network to control a GAN model enables an out-of-representation input image, which cannot be accurately represented by a GAN model, to be successfully mapped to the valid representation space of the GAN. In this paper, we exploit a hyper-network for the few-shot multi-domain adaptation problem.

### 3 DYNAGAN

We aim to adapt a GAN model pretrained on a source domain to multiple target domains with only a few target images, e.g., from the real-portrait FFHQ dataset [Karras et al. 2019] to cartoon, anime, and painting with only one or two images for each. To this end, we set three design goals for our method. First, our model should be capable of synthesizing diverse images reflecting distinct target domain characteristics, not an average style of them. Second, our model should be compact in model size in order to provide benefits over using a separate model for each target domain. Lastly, our model should be controllable in choosing which target domain to use for image synthesis. In the following, we present our framework DynaGAN and its training strategy to meet these requirements.

#### 3.1 Framework

DynaGAN consists of a generator network and an adaptation module as shown in Fig. 2. For the generator network, we adopt a StyleGAN2 generator [Karras et al. 2020b] pretrained on a source domain. The generator synthesizes an image from a GAN latent

vector  $z$ , and the style of the synthesized image is controlled by a target domain condition  $c$  through the adaptation module. Here, the adaptation module acts as a hyper-network that dynamically modulates the weights of the generator according to a given target domain condition  $c$ . As we achieve domain adaptation using the adaptation module while keeping the generator intact, the generator can still synthesize images of the original source domain by turning off modulation from the adaptation module.

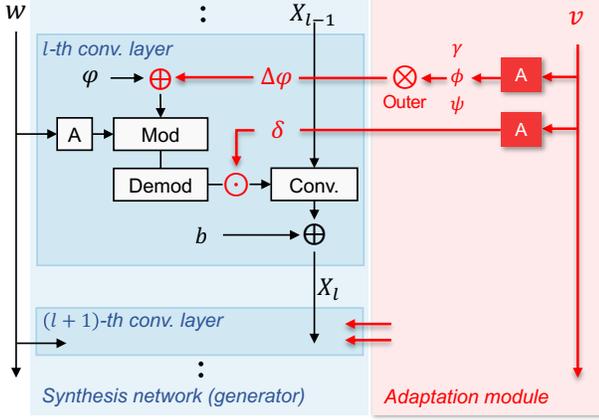
*Modulation-aware Generator.* For the generator network, we adopt the StyleGAN2 generator architecture, which consists of a mapping network and a synthesis network, with simple architectural changes, which will be described in the next paragraph. Here, we first describe the original StyleGAN2 generator. As illustrated in Fig. 2, the mapping network  $M_G$  converts a latent vector  $z$  to an intermediate latent vector  $w$ , which is fed to the convolutional layers of a synthesis network. As shown in Fig. 3, a synthesis network has multiple convolutional layers. We denote the filter weights of the  $l$ -th convolutional layer as  $\varphi \in \mathbb{R}^{C_{in} \times C_{out} \times k}$  where  $C_{in}$  is the input channel size,  $C_{out}$  is the number of output filters, and  $k$  is the spatial size of the filter. In the original StyleGAN2 architecture, the  $l$ -th convolutional layer transforms an input feature  $X_{l-1}$  from the previous layer into an output feature  $X_l$  as:

$$X_l = X_{l-1} * f(\varphi, A(w)) + b, \quad (1)$$

where  $*$  is the convolution operator and  $b$  is the bias of the  $l$ -th convolutional layer.  $f(\cdot)$  is a composited function defined as  $f = \text{Demod} \circ \text{Mod}(\cdot)$ , where  $\text{Mod}(\cdot)$  and  $\text{Demod}(\cdot)$  are the modulation and demodulation operations, respectively.  $A(\cdot)$  is an affine transformation layer. For detailed definitions of the operations and layers, we refer the readers to StyleGAN2 [Karras et al. 2020b].

In order to provide effective modulation on the generator, we change the architecture of the convolutional layer by introducing additional modulation parameters as:

$$X_l = X_{l-1} * (\delta \cdot f(\varphi + \Delta\varphi, A(w))) + b, \quad (2)$$



**Figure 3: Parameter modulation of the generator by the adaptation module.**

where  $\delta \in \mathbb{R}^{C_{\text{out}}}$  is a filter-wise scaling factor applied to the output of the function  $f$  as  $(\delta \cdot \varphi)_{ijk} = \delta_j \varphi_{ijk}$ , where  $i, j$  and  $k$  are indices of the input channel, output channel and spatial location, respectively.  $\Delta\varphi$  is a residual weight added to the filter weight. The two modulation parameters,  $\delta$  and  $\Delta\varphi$ , basically shift and scale the weights of the convolutional layer. Note that we modulate the convolution weights without touching the other parts of the generator, e.g. the mapping network. This is aligned with the observations made in recent methods that changing convolutional weights is effective for domain adaptation [Pinkney and Adler 2020; Wu et al. 2021, 2022]. Our approach also uses filter-wise scaling instead of a general linear transformation for  $\varphi$ , which is motivated by the observation of Wu et al. [Wu et al. 2021] that images can be successfully modified by scaling intermediate feature maps of a generator in a channel-wise manner. This filter-wise scaling significantly reduces the number of parameters, while still achieving effective domain adaptation. We highlight our modification as red text in Fig. 3.

*Efficient Adaptation Module.* We then introduce our adaptation module that estimates the modulation parameters  $\Delta\varphi$  and  $\delta$ . See Fig. 3 for its network architecture. Given a target domain one-hot vector  $c$ , we transform the vector  $c$  into a latent vector  $v$  using a multi-layer perceptron (MLP)-based mapping network. The latent vector  $v$  is then fed to the affine layers of the adaptation module, resulting in the modulation parameters  $\Delta\varphi$  and  $\delta$ . Each affine layer is composed of a single fully-connected layer.

Estimating the residual parameters  $\Delta\varphi$  for all the layers poses a significant burden in memory and other training resources, as the StyleGAN2 generator has nearly 2.3M parameters per convolutional layer. To tackle this issue, we represent the residual parameter  $\Delta\varphi$  with its rank-1 tensor decomposition [Hou and Kung 2021] as

$$\Delta\varphi = \gamma \otimes \phi \otimes \psi, \quad (3)$$

where  $\gamma \in \mathbb{R}^{C_{\text{out}}}$ ,  $\phi \in \mathbb{R}^{C_{\text{in}}}$  and  $\psi \in \mathbb{R}^k$  are decomposed 1D vectors. Using rank-1 decomposed vectors significantly reduces the burden of our adaptation network by only estimating three 1D vectors, instead of a high-dimensional tensor  $\Delta\varphi$ . Our rank-decomposed representation is more efficient than channel-wise residuals [Alaluf et al. 2022]. Refer to the Supplemental Document for details.

## 3.2 Training

*Initialization.* We train our adaptation module to estimate the rank-1 decomposed modulation parameters  $\gamma$ ,  $\phi$ ,  $\psi$  and the filter-wise scaling factor  $\delta$ . The generator network is fixed during training the adaptation module. For a warm start of the training process, we initialize our adaptation module so that the modulation parameters do not significantly change the original weights in the pretrained generator network. Specifically, we initialize the weights of the affine layers of the adaptation module with random initialization attenuated by a linear scalar 0.01. Furthermore, the bias parameter of the  $\delta$ -estimation affine layer is set to one to make  $\delta$  close to one. For  $\gamma$ ,  $\phi$  and  $\psi$ , the biases of the corresponding affine layers are initialized to 1, 1, and 0 respectively, resulting in  $\Delta\varphi$  close to zero.

*Multi-domain Adaptation Loss.* For few-shot multi-domain adaptation, we design our training loss as a weighted sum of three loss functions:

$$\mathcal{L} = \lambda_{\text{contra}} \mathcal{L}_{\text{contra}} + \lambda_{\text{MTG}} \mathcal{L}_{\text{MTG}} + \lambda_{\text{ID}} \mathcal{L}_{\text{ID}}, \quad (4)$$

where  $\lambda_{\text{contra}}$ ,  $\lambda_{\text{MTG}}$  and  $\lambda_{\text{ID}}$  are the balancing weights.  $\mathcal{L}_{\text{contra}}$ ,  $\mathcal{L}_{\text{MTG}}$  and  $\mathcal{L}_{\text{ID}}$  are contrastive-adaptation loss, MTG loss and an identity loss, respectively. In our experiments, we set  $\lambda_{\text{contra}} = 1$ ,  $\lambda_{\text{MTG}} = 1$  and  $\lambda_{\text{ID}} = 3$  if both source and target domains contain faces, and  $\lambda_{\text{ID}} = 0$  otherwise.

First, our novel contrastive-adaptation loss  $\mathcal{L}_{\text{contra}}$  promotes keeping the distinct characteristics of different target domains. The key idea behind the contrastive-adaptation loss is to make the synthesized image be similar/different with a target training image, for positive/negative pairs, which means that they are in the same/different target domains. We denote a training image and a synthesized image at a target domain  $c$  as  $I_c$  and  $\hat{I}_c(w)$ , respectively, where  $w$  is a StyleGAN2 latent vector. We evaluate the similarity between the two images using cosine similarity in a recently-proposed CLIP embedding space [Radford et al. 2021]. Specifically, the contrastive-adaptation loss  $\mathcal{L}_{\text{contra}}$  is formulated as

$$\mathcal{L}_{\text{contra}} = -\log \frac{\exp(I_{\text{pos}}/\tau)}{\exp(I_{\text{pos}}/\tau) + \sum_j \mathbb{1}_{[j \neq c]} \exp(I_{\text{neg}}^j/\tau)}, \quad (5)$$

where  $\tau$  is a temperature parameter, which is set to one.  $I_{\text{pos}}$  and  $I_{\text{neg}}^j$  are similarities of positive and negative pairs defined as:

$$\begin{aligned} I_{\text{pos}} &= \text{sim}(E_{\text{CLIP}}(I_c), E_{\text{CLIP}}(\hat{I}_c(w))), \\ I_{\text{neg}}^j &= \text{sim}(E_{\text{CLIP}}(\text{Aug}(I_j)), E_{\text{CLIP}}(\hat{I}_c(w))), \end{aligned} \quad (6)$$

where  $\text{sim}(\cdot)$  is cosine similarity and  $E_{\text{CLIP}}$  is a CLIP encoder [Radford et al. 2021].  $\text{Aug}(\cdot)$  is a horizontal-flip and color-jitter augmentation function for stable training [Liu et al. 2021].

For  $\mathcal{L}_{\text{MTG}}$ , we adopt the training loss of MTG [Zhu et al. 2022], which is defined as a combination of a reconstruction loss and CLIP-based losses. To train our adaptation module, we made a few modifications to the original MTG loss to inject domain dependency, which is detailed in the Supplemental Document. The identity loss  $\mathcal{L}_{\text{ID}}$  measures the similarity of faces between synthesized and source-domain images in the representation space of a face recognition network [Deng et al. 2019]. Also, refer to the Supplemental Document for its details.

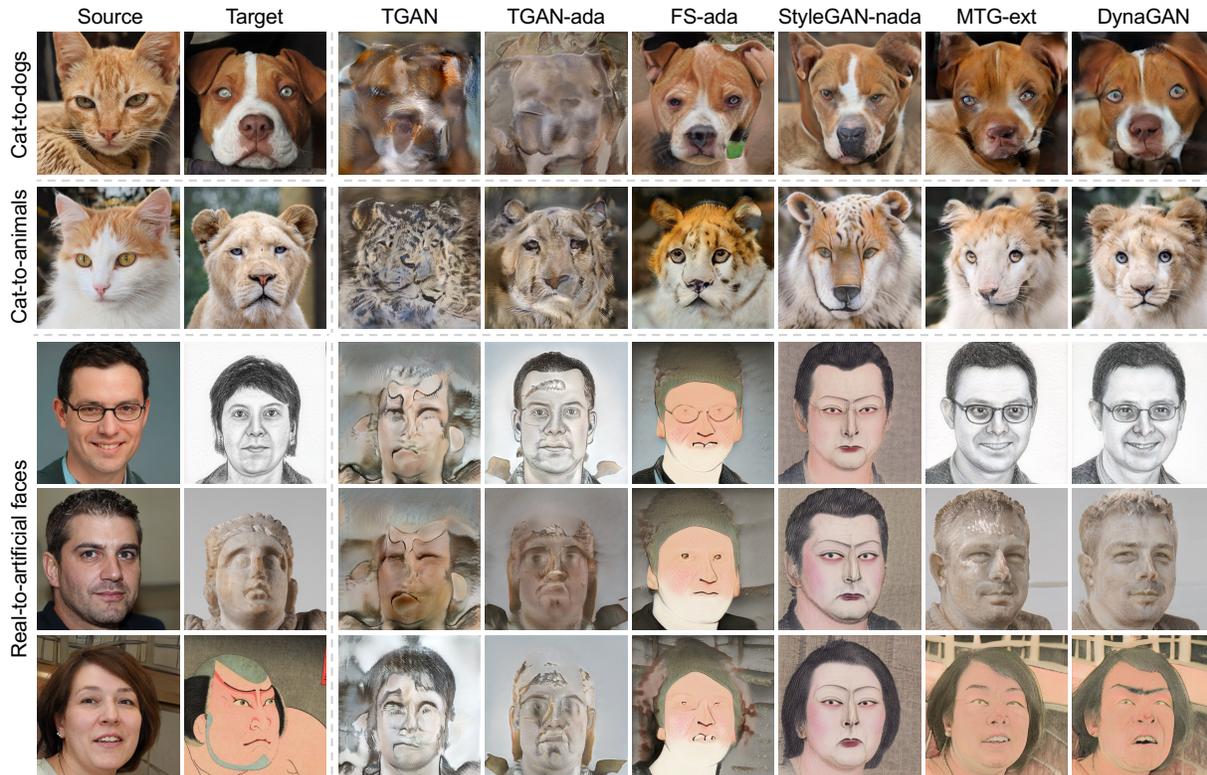


Figure 4: Qualitative comparison against state-of-the-art few-shot domain adaptation methods: TGAN [Wang et al. 2018], TGAN-ada [Karras et al. 2020a], FS-ada [Ojha et al. 2021], StyleGAN-nada [Gal et al. 2022], and MTG-ext [Zhu et al. 2022]. Targets in the 1st, 2nd, 4th and 5th rows: Pixabay (Doz777, PhotoGranary) [Pixabay License], The Metropolitan Museum of Art [Public Domain], and ARC Collection, Ritsumeikan University (arcBK01-0042\_37).

## 4 EXPERIMENTS

### 4.1 Implementation Details

For training the adaptation module, we use the Adam optimizer [Kingma and Ba 2015] with coefficients  $\beta_1 = 0.0$  and  $\beta_2 = 0.99$ . We set the learning rate to 0.002 and the batch size to 4. When we compute the loss functions, we should embed an image into a CLIP representation space, for which we use ‘ViT-B/16’ and ‘ViT-B/32’ CLIP encoder models [Radford et al. 2021] and add their results as done in MTG [Zhu et al. 2022].

At the inference time, we apply the style mixing of MTG [Zhu et al. 2022] to our generator shown in Fig. 2 in order to better reflect the style of a target domain. Specifically, for *fine-scale* layers of our synthesis network, we use a latent vector  $w_c$  estimated from a target-domain training image rather than using a latent vector  $w$  from the mapping network of the generator  $M_G$  to more faithfully transfer texture styles of the target domain. For the rest of the layers, we use the latent vector  $w$ . More details on the implementation can be found in the Supplemental Document.

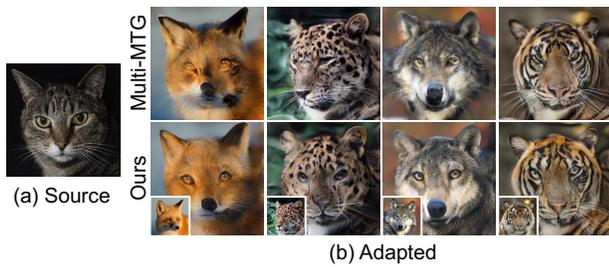
### 4.2 Comparison

*Comparison Methods.* As DynaGAN is the first few-shot multi-domain adaptation method, we instead compare our method with five state-of-the-art single-domain adaptation methods: TGAN [Wang et al. 2018], TGAN-ada [Karras et al. 2020a], FS-ada [Ojha et al.

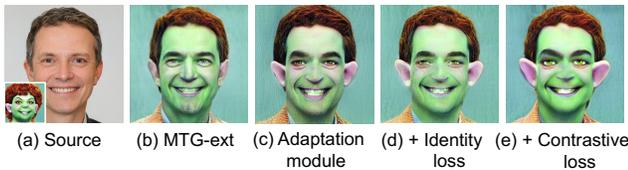
2021], StyleGAN-nada [Gal et al. 2022], and MTG-ext, which is our extension of MTG [Zhu et al. 2022]. While the original MTG is a single-shot adaptation method, we extend it to handle few-shot training images by slightly modifying its loss and to use a particular training image for its style-mixing step. More details on MTG-ext can be found in the Supplemental Document.

*Experimental Datasets.* To evaluate the effectiveness of DynaGAN compared to the state-of-the-art methods, we curate three challenging few-shot multi-domain datasets: (a) cat-to-dogs, (b) cat-to-animals, and (c) real-to-artificial faces. For the three datasets, each image corresponds to an *individual* target domain, meaning that each target domain has only one sample image. Refer to the Supplemental Document for the experiments that use multiple samples per domain. For (a) the cat-to-dogs dataset, we use five target images sampled from the AFHQ Dog dataset [Choi et al. 2020] and the source domain is AFHQ Cat dataset [Choi et al. 2020]. In (b) the cat-to-animals dataset, we have 10 target images of different animal species from the AFHQ Wild dataset [Choi et al. 2020]. AFHQ Cat dataset [Choi et al. 2020] is used again for the source domain. For (c) the real-to-artificial faces, we use nine target images with diverse forms consisting of three sketch images<sup>1</sup>, three samples from the MetFace dataset [Karras et al. 2020a], and three Ukiyo-e

<sup>1</sup>We synthesized sketch images from the Sketch dataset using FS-ada [Ojha et al. 2021] instead of directly using the dataset.



**Figure 5: Qualitative comparison with multi-MTGs on the cat-to-animals dataset. Multi-MTGs consists of multiple MTG [Zhu et al. 2022] models separately trained on each target domain. Despite using a single model, our method represents each domain more effectively. Insets from left to right: Pixabay (Pexels, vnzling, WikiImages, Gregorius\_o) [Pixabay License].**

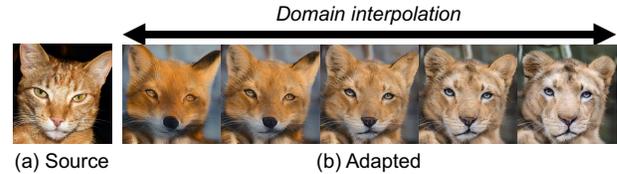


**Figure 6: For (a) the source image, (b) our baseline model (MTG-ext) fails to synthesize an image that faithfully reflects the style of the target domain. (c) Our adaptation module allows us to reflect the style of the target domain by means of dynamic modulation. (d) Identity loss helps preserving the face identity. (e) Our contrastive-adaptation loss further helps reflect domain-specific attributes.**

images [Pinkney 2020]. We use the FFHQ [Karras et al. 2019] dataset as the source domain.

*Qualitative Assessment.* Fig. 4 shows synthesized images using each of the adapted models on the three target datasets. TGAN [Wang et al. 2018] and TGAN-ada [Karras et al. 2020a] fail to learn faithful distributions of the target-domain images, resulting in low-fidelity synthesis. FS-ada [Ojha et al. 2021] and StyleGAN-nada [Gal et al. 2022] produce images corresponding to average styles of target domains, losing the distinct characteristics of different target domains. MTG-ext achieves better results with more faithful textures than the previous methods thanks to the style mixing technique. However, it suffers from unnatural global structures, e.g., the deformed dog head in the first row of Fig. 4. DynaGAN achieves the high-quality images via our adaptation module that dynamically adapts the generator.

*Quantitative Assessment.* We perform quantitative evaluation using Fréchet inception distance (FID) [Heusel et al. 2017], inception score (IS) [Salimans et al. 2016] and kernel inception distance (KID) [Bińkowski et al. 2018]. For the real target-domain images, we use 171 dog images, 483 animal images, and 900 face images in the three datasets. We use 50,000 synthesized images for each method. Tab. 1 shows that DynaGAN achieves superior results to all the other methods, especially for target datasets with large variations: cat-to-animals dataset and real-to-artificial faces dataset.



**Figure 7: DynaGAN supports domain interpolation. For (a) the source image of a cat, we synthesize images with different target domains, here from fox to lion.**

*Separate Models for Multi-domain Adaptation.* As mentioned in Sec. 1, a naïve approach to multi-domain adaptation is to train a separate model, here we used MTG [Zhu et al. 2022], for each target domain, called multi-MTGs. To compare DynaGAN against the naïve approach, we train a separate MTG model for each target domain in the cat-to-animals dataset. Note that with 10 different target domains in the cat-to-animals dataset, multi-MTGs result in  $10\times$  network parameters compared to the original MTG [Zhu et al. 2022], resulting in 302.8M parameters. Note that DynaGAN has only 39.6M parameters, being more light-weight. Fig. 5 shows a qualitative comparison between DynaGAN and multi-MTGs. DynaGAN produces higher-quality results better reflecting the distinct attributes of the target domains using a single generator and an adaptation module.

### 4.3 Ablation Study

Fig. 6 shows an ablation study where we add the components of DynaGAN one by one to our baseline model, which does not use our adaptation module and the contrastive-adaptation loss. Here, the target dataset is curated with 5 different images with diverse styles sampled from StyleGAN-nada [Gal et al. 2022]. Our adaptation module significantly improves the quality of synthesized images with more plausible shapes of human faces. Our contrastive-adaptation loss leads to a result that better reflects domain-specific attributes.

### 4.4 Applications

Here, we demonstrate applications of DynaGAN: domain interpolation, image-to-image translation, adaptation degree control, and semantic image manipulation.

*Domain Interpolation.* Our adaptation module provides a smooth transition between two target domains by interpolating the target-domain vector  $c$ . Fig. 7 shows an example of domain interpolation where the source domain is cat, and we interpolate the target-domain vector from fox to lion.

*Image-to-image Translation.* Image-to-image translation is a task that transfers an image of a source domain to a target domain. By combining DynaGAN with GAN inversion, we demonstrate successful image-to-image translation of real-world images. Fig. 8 shows results of image-to-image translation. We invert the input images to GAN latent codes using an off-the-shelf GAN inversion method [Tov et al. 2021], and synthesize multi-domain images.

*Controllable Degree of Adaptation.* DynaGAN allows to control the degree of adaptation at the inference time by adjusting (1) the modulation parameters  $\Delta\phi$  and  $\delta$  in Eq. (2) and (2) the degree of style mixing. Specifically, we can scale the modulation parameters

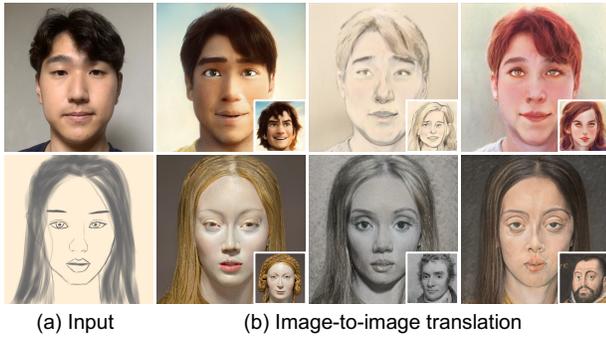


Figure 8: Image-to-image translation of real-world images. The inset images in the translation results are training samples of the target domains. 1st, 2nd and 3rd insets in the 2nd row: The Metropolitan Museum of Art [Public Domain].

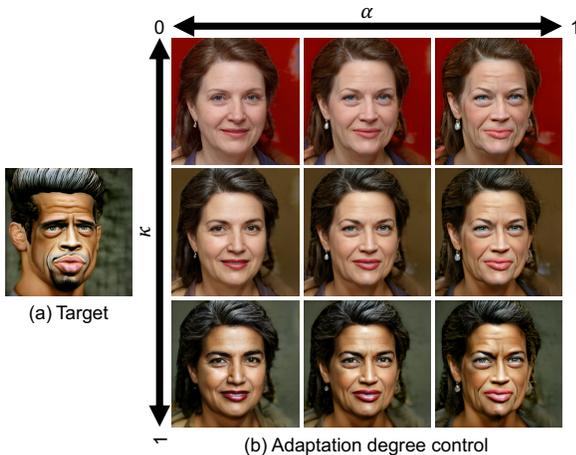


Figure 9: Adaptation degree control using the adaptation control parameter  $\alpha$  and the style mixing parameter  $\kappa$ .

as  $\Delta\varphi \leftarrow \alpha\Delta\varphi$  and  $\delta \leftarrow \alpha\delta + (1 - \alpha)$ , where  $\alpha$  is the control parameter that adjusts the degree of adaptation. For the degree of style mixing, we compute an *interpolated* latent vector  $\hat{w}$  between  $w$  and  $w_c$  as  $\hat{w} = (1 - \kappa)w + \kappa(w_c)$  where  $\kappa$  is the control parameter. We feed  $\hat{w}$  instead of  $w_c$  to fine-scale convolutional layers of the generator. Fig. 9 shows results with varying control parameters of  $\alpha$  and  $\kappa$  from 0 to 1.  $\alpha$  controls the adaptation degree of shape and  $\kappa$  takes care of texture.

*Semantic Image Manipulation.* As DynaGAN keeps the mapping network of the generator  $M_G$  intact, source and target domains share the common latent space [Wu et al. 2022]. This allows manipulation of target-domain images in a GAN latent space in a same way in source-domain images. Fig. 10 shows an example of semantic image manipulation of a target-domain image where the image is manipulated using StyleCLIP [Patashnik et al. 2021] and InterfaceGAN [Shen et al. 2020].

## 5 CONCLUSION

In this paper, we propose DynaGAN, which effectively adapts a pretrained generator to diverse target domains with just a few target images. Our adaptation module is light-weight and provides

Table 1: Quantitative comparison of different methods on different dataset scenarios.

Model	Metric	Dataset		
		Cat-to-dogs	Cat-to-animals	Real-to-artificial faces
TGAN	FID ↓	342.97	179.96	182.42
	IS ↑	<b>2.40</b>	1.66	2.03
	KID $\times 10^3$ ↓	411.11	160.90	106.64
TGAN+ada	FID ↓	238.57	175.46	178.09
	IS ↑	<b>1.61</b>	1.88	3.11
	KID $\times 10^3$ ↓	285.13	141.79	118.60
FS-ada	FID ↓	73.29	<b>80.80</b>	148.13
	IS ↑	1.31	1.92	1.92
	KID $\times 10^3$ ↓	52.13	<b>53.72</b>	109.31
StyleGAN-nada	FID ↓	60.77	108.50	142.41
	IS ↑	1.07	1.18	1.77
	KID $\times 10^3$ ↓	35.61	84.57	102.28
MTG-ext	FID ↓	<b>55.84</b>	89.21	<b>104.59</b>
	IS ↑	1.24	<b>2.37</b>	<b>4.09</b>
	KID $\times 10^3$ ↓	<b>29.60</b>	58.48	48.77
DynaGAN	FID ↓	<b>55.08</b>	<b>38.37</b>	<b>97.23</b>
	IS ↑	1.34	<b>4.53</b>	<b>4.03</b>
	KID $\times 10^3$ ↓	<b>23.86</b>	<b>17.36</b>	<b>41.35</b>

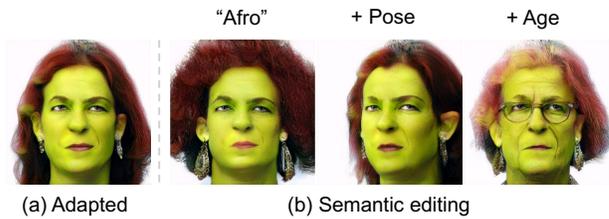


Figure 10: Semantic editing example. (a) A target-domain image. (b) We add "Afro" hair using the text-driven editing of StyleCLIP [Patashnik et al. 2021], and change the pose and the age using InterFaceGAN [Shen et al. 2020].



Figure 11: Failure cases of DynaGAN. (a) and (c) are target-domain training images, while (b) and (d) are synthesized results by adapted generators. The inset images in (a) and (c) are source-domain images generated from the same latent codes used for (b) and (d). Target in (a): Pixabay (Joe007) [Pixabay License].

sufficient expressive power for the generator, allowing us to handle multiple target domains. Our contrastive loss further encourages generating distinct attributes of different domains. We demonstrate that DynaGAN can effectively handle multiple target domains and achieve superior results to previous state-of-the-art methods.

**Limitations.** DynaGAN fails when there exists a large domain gap between the source and target domains, such as human faces and dogs, as shown in Fig. 11(a) and (b). This is because DynaGAN fundamentally depends on the source-domain knowledge in the pretrained network for domain adaptation. Also, using extremely few-shot training images may raise overfitting. For example, in Fig. 11(c) and (d), even though its source-domain image has no glasses, the adaptation result has thin structures like glasses around the eyes as the target image has glasses.

## ACKNOWLEDGMENTS

This research was supported by IITP grants funded by the Korea government (MSIT) (2021-0-02068, 2019-0-01906), an NRF grant funded by the the Korea government (MOE) (2022R1A6A1A03052954), and Peblblous.

## REFERENCES

- Yuval Alaluf, Omer Tov, Ron Mokady, Rinon Gal, and Amit Bermano. 2022. Hyperstyle: Stylegan inversion with hypernetworks for real image editing. In *Proc. of IEEE/CVF CVPR*.
- Mikołaj Bińkowski, Dougal J. Sutherland, Michael Arbel, and Arthur Gretton. 2018. Demystifying MMD GANs. In *Proc. of ICLR*.
- Andrew Brock, Jeff Donahue, and Karen Simonyan. 2019. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *Proc. of ICLR*.
- Kelvin CK Chan, Xintao Wang, Xiangyu Xu, Jinwei Gu, and Chen Change Loy. 2021. GLEAN: Generative Latent Bank for Large-Factor Image Super-Resolution. In *Proc. of IEEE/CVF CVPR*.
- Yunjeong Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. 2020. StarGAN v2: Diverse Image Synthesis for Multiple Domains. In *Proc. of IEEE/CVF CVPR*.
- Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. 2019. ArcFace: Additive angular margin loss for deep face recognition. In *Proc. of IEEE/CVF CVPR*.
- Tan M Dinh, Anh Tuan Tran, Rang Nguyen, and Binh-Son Hua. 2022. Hyperinverter: Improving stylegan inversion via hypernetwork. In *Proc. of IEEE/CVF CVPR*.
- Rinon Gal, Or Patashnik, Haggai Maron, Amit H. Bermano, Gal Chechik, and Daniel Cohen-Or. 2022. StyleGAN-NADA: CLIP-Guided Domain Adaptation of Image Generators. *ACM Trans. Graph.* 41, 4, Article 141 (jul 2022), 13 pages. <https://doi.org/10.1145/3528223.3530164>
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Proc. of NeurIPS*.
- David Ha, Andrew M. Dai, and Quoc V. Le. 2017. HyperNetworks. In *Proc. of ICLR*.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Proc. of NeurIPS*.
- Zejiang Hou and Sun-Yuan Kung. 2021. Parameter Efficient Dynamic Convolution via Tensor Decomposition. In *Proc. of BMVC*.
- Jialu Huang, Jing Liao, and Sam Kwong. 2022. Unsupervised Image-to-Image Translation via Pre-Trained StyleGAN2 Network. *IEEE Transactions on Multimedia* 24 (2022), 1435–1448. <https://doi.org/10.1109/TMM.2021.3065230>
- Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. 2021. Deceive D: Adaptive Pseudo Augmentation for GAN Training with Limited Data. In *Proc. of NeurIPS*.
- Kyoungkook Kang, Seongtae Kim, and Sunghyun Cho. 2021. GAN Inversion for Out-of-Range Images with Geometric Transformations. In *Proc. of IEEE/CVF ICCV*.
- Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. 2020a. Training generative adversarial networks with limited data. In *Proc. of NeurIPS*.
- Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2021. Alias-Free Generative Adversarial Networks. In *Proc. of NeurIPS*.
- Tero Karras, Samuli Laine, and Timo Aila. 2019. A style-based generator architecture for generative adversarial networks. In *Proc. of IEEE/CVF CVPR*.
- Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. 2020b. Analyzing and improving the image quality of StyleGAN. In *Proc. of IEEE/CVF CVPR*.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proc. of ICLR*.
- Yijun Li, Richard Zhang, Jingwan Lu, and Eli Shechtman. 2020. Few-shot image generation with elastic weight consolidation. (2020).
- Gidi Littwin and Lior Wolf. 2019. Deep meta functionals for shape representation. In *Proc. of IEEE/CVF ICCV*.
- Xingchao Liu, Chengyue Gong, Lemeng Wu, Shujian Zhang, Hao Su, and Qiang Liu. 2021. Fusedream: Training-free text-to-image generation with improved clip+ gan space optimization. *arXiv preprint arXiv:2112.01573* (2021).
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2015. Deep learning face attributes in the wild. In *Proc. of IEEE/CVF ICCV*.
- Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, and James Henderson. 2021. Parameter-efficient multi-task fine-tuning for transformers via shared hypernetworks. *arXiv preprint arXiv:2106.04489* (2021).
- Yuval Nirkin, Lior Wolf, and Tal Hassner. 2021. HyperSeg: Patch-wise hypernetwork for real-time semantic segmentation. In *Proc. of IEEE/CVF CVPR*.
- Atsuhiko Noguchi and Tatsuya Harada. 2019. Image generation from small datasets via batch statistics adaptation. In *Proc. of IEEE/CVF ICCV*.
- Utkarsh Ojha, Yijun Li, Jingwan Lu, Alexei A Efros, Yong Jae Lee, Eli Shechtman, and Richard Zhang. 2021. Few-shot image generation via cross-domain correspondence. In *Proc. of IEEE/CVF CVPR*.
- Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. 2021. StyleCLIP: Text-Driven Manipulation of StyleGAN Imagery. In *Proc. of IEEE/CVF ICCV*.
- Justin NM Pinkney and Doron Adler. 2020. Resolution dependent GAN interpolation for controllable image synthesis between domains. *arXiv preprint arXiv:2010.05334* (2020).
- Justin N. M. Pinkney. 2020. Aligned Ukiyo-e faces dataset. <https://www.justinpinkney.com/ukiyoe-dataset>.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *Proc. of ICLR*.
- Neale Ratzlaff and Li Fuxin. 2019. HyperGAN: A generative model for diverse, performant neural networks. In *Proc. of ICLR*.
- Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. 2021. Encoding in Style: a StyleGAN Encoder for Image-to-Image Translation. In *Proc. of IEEE/CVF CVPR*.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. 2016. Improved techniques for training GANs. In *Proc. of NeurIPS*.
- Veit Sandfort, Ke Yan, Perry J Pickhardt, and Ronald M Summers. 2019. Data augmentation using generative adversarial networks (CycleGAN) to improve generalizability in CT segmentation tasks. *Scientific reports* (2019).
- Yujun Shen, Jinjin Gu, Xiaou Tang, and Bolei Zhou. 2020. Interpreting the Latent Space of GANs for Semantic Face Editing. In *Proc. of IEEE/CVF CVPR*.
- Yujun Shen and Bolei Zhou. 2021. Closed-Form Factorization of Latent Semantics in GANs. In *Proc. of IEEE/CVF CVPR*.
- Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. 2021. Designing an Encoder for StyleGAN Image Manipulation. *ACM Trans. Graph.* 40, 4, Article 133 (jul 2021), 14 pages. <https://doi.org/10.1145/3450626.3459838>
- Rotem Tzaban, Ron Mokady, Rinon Gal, Amit H Bermano, and Daniel Cohen-Or. 2022. Stitch it in Time: GAN-Based Facial Editing of Real Videos. *arXiv preprint arXiv:2201.08361* (2022).
- Johannes von Oswald, Christian Henning, Benjamin F. Grewe, and João Sacramento. 2020. Continual learning with hypernetworks. In *Proc. of ICLR*.
- Xintao Wang, Yu Li, Honglun Zhang, and Ying Shan. 2021. Towards real-world blind face restoration with generative facial prior. In *Proc. of IEEE/CVF CVPR*.
- Yaxing Wang, Abel Gonzalez-Garcia, David Berga, Luis Herranz, Fahad Shahbaz Khan, and Joost van de Weijer. 2020. MineGAN: effective knowledge transfer from GANs to target domains with few images. In *Proc. of IEEE/CVF CVPR*.
- Yaxing Wang, Chenshen Wu, Luis Herranz, Joost van de Weijer, Abel Gonzalez-Garcia, and Bogdan Raducanu. 2018. Transferring GANs: generating images from limited data. In *Proc. of ECCV*.
- Zongze Wu, Dani Lischinski, and Eli Shechtman. 2021. StyleSpace analysis: Disentangled controls for StyleGAN image generation. In *Proc. of IEEE/CVF CVPR*.
- Zongze Wu, Yotam Nitzan, Eli Shechtman, and Dani Lischinski. 2022. StyleAlign: Analysis and Applications of Aligned StyleGAN Models. In *Proc. of ICLR*.
- Tao Yang, Peiran Ren, Xuansong Xie, and Lei Zhang. 2021. GAN prior embedded network for blind face restoration in the wild. In *Proc. of IEEE/CVF CVPR*.
- Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. 2015. LSUN: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365* (2015).
- Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. 2020. Differentiable augmentation for data-efficient GAN training. In *Proc. of NeurIPS*.
- Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. 2020. In-domain GAN Inversion for Real Image Editing. In *Proc. of ECCV*.
- Peihao Zhu, Rameen Abdal, John Femiani, and Peter Wonka. 2022. Mind the Gap: Domain Gap Control for Single Shot Domain Adaptation for Generative Adversarial Networks. In *Proc. of ICLR*.
- Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. 2018. Learning transferable architectures for scalable image recognition. In *Proc. of IEEE/CVF CVPR*.