

Polymorphic-GAN: Generating Aligned Samples across Multiple Domains with Learned Morph Maps

Seung Wook Kim^{1,2,3} Karsten Kreis¹ Daiqing Li¹

Antonio Torralba⁴ Sanja Fidler^{1,2,3}

¹NVIDIA ²University of Toronto ³Vector Institute ⁴MIT

{seungwookk, kkreis, daiqingl, sfidler}@nvidia.com, torralba@mit.edu

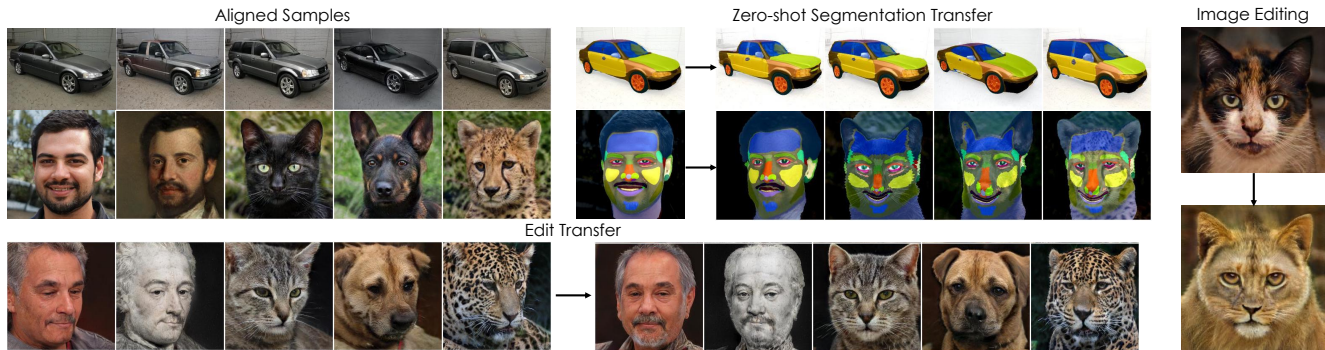


Figure 1. We introduce *Polymorphic-GAN* (PMGAN) for generating aligned samples across multiple domains. PMGAN enables a diverse set of applications, including zero-shot segmentation and cross-domain image editing, by learning geometric differences between domains.

Abstract

Modern image generative models show remarkable sample quality when trained on a single domain or class of objects. In this work, we introduce a generative adversarial network that can simultaneously generate aligned image samples from multiple related domains. We leverage the fact that a variety of object classes share common attributes, with certain geometric differences. We propose *Polymorphic-GAN* which learns shared features across all domains and a per-domain morph layer to morph shared features according to each domain. In contrast to previous works, our framework allows simultaneous modelling of images with highly varying geometries, such as images of human faces, painted and artistic faces, as well as multiple different animal faces. We demonstrate that our model produces aligned samples for all domains and show how it can be used for applications such as segmentation transfer and cross-domain image editing, as well as training in low-data regimes. Additionally, we apply our *Polymorphic-GAN* on image-to-image translation tasks and show that we can greatly surpass previous approaches in cases where the geometric differences between domains are large.

1. Introduction

Generative adversarial networks (GANs) have achieved remarkable image synthesis quality [5, 11, 27, 28]. Moreover, GANs like StyleGAN [27, 28] have been shown to form a semantic understanding of the modeled images in

their features [3, 4, 14, 21, 24, 34, 54, 59, 61, 66, 70], which has been leveraged in diverse applications, including image editing [9, 19, 30, 31, 36, 58, 73], inverse rendering [68], style transfer [1, 29], image-to-image translation [7, 8, 22, 49], and semi-supervised learning [34, 66, 70].

GANs are usually trained on images from individual domains, such as human faces [25, 27]. However, there are many related domains which share similar semantics and characteristics, such as animal faces or face paintings. In our work, we aim to train a generative model with a shared backbone to produce *aligned* samples from multiple related domains. By aligned, we mean images that share common attributes and conditions across domains, such as pose and lighting. This has an obvious computational advantage by sharing weights across domains, but more importantly, it affords a variety of applications such as transferring segmentation labels from one domain to another, in which such information may not be available. Furthermore, by editing one domain, we get edits in other domains for free.

The main obstacle to building a GAN that simultaneously synthesizes outputs from different domains is that even though the *semantics* are often shared, the *geometry* can vary significantly (consider, for example, the face of a human, a dog, and a cat). This prevents a natural sharing of generator features among such semantically aligned, but geometrically varying domains. Common approaches such as fine-tuning a pre-trained GAN [26, 43, 64] unfortunately lose the ability to sample from the parent domain, or, more generally, multiple domains at the same time. Learn-

ing shared representations between multiple domains has been studied in the transfer and multi-task learning literature [10, 39, 40, 72], but there has been little progress in generative models [2, 8, 20, 37, 38].

To overcome these challenges, we propose *Polymorphic-GAN* (PMGAN). It leverages a shared generator network together with novel *morph maps* that geometrically deform and adapt the synthesis network’s feature maps. In particular, PMGAN builds on the StyleGAN2 [28] architecture and augments the model with a *MorphNet* that predicts domain-specific morph maps which warp the main generator’s features according to the different domains’ geometries. An additional shallow convolutional neural network is then sufficient to render these morphed features into correctly stylized and geometrically aligned outputs that are also semantically consistent across multiple domains.

By sharing as many generator layers as possible, the impressive semantic properties of StyleGAN’s latent space are shared across all modeled domains, while the geometric differences are still correctly reflected due to the additional morph operations. Because of that, our PMGAN enables many relevant applications in a unique and novel way. We extensively analyze PMGAN and validate the method on the following tasks: (i) We perform previously impossible expressive image editing across different domains. (ii) We use PMGAN for image-to-image translation across domains, and outperform previous methods in cases where the geometric gap between domains is large. (iii) We leverage PMGAN’s learnt morph maps for zero-shot semantic segmentation transfer across domains. (iv) Finally, sharing the generator’s features across domains is advantageous when involving domains with little training data. In these cases the main generator can be learnt primarily from a domain with much data, which benefits all other domains. In summary, our PMGAN is the first generative model that naturally and easily allows users to synthesize aligned samples from multiple semantically-related domains at the same time, enabling novel and promising applications.

2. Related Work

StyleGAN. StyleGAN [27, 28] is the state-of-the-art GAN model with remarkable sample quality, which has enabled many relevant applications. StyleGAN inversion methods [1, 49, 58, 73] discover the latent vector corresponding to an input image. Once an image is embedded, GAN-based editing methods [36, 53, 54, 64] find semantically meaningful directions in latent space to achieve desired editing effects. DatasetGAN [70] and SematicGAN [34] use StyleGAN’s feature maps for producing segmentation masks.

StyleGAN Adaptation. Various methods adapt a pre-trained StyleGAN for a target domain. Fine-tuning approaches [26, 43, 64] take a pre-trained model as a starting point for optimization and learn to generate samples from

the target domain. Few-shot adaptation approaches [12, 45, 47] make use of a small amount of data from the target domain or CLIP [48] to adapt latent codes or model weights towards the target domain. PMGAN, in contrast, learns a single generator that jointly models multiple domains.

Cross-Domain Generation. Several works [7, 8, 20, 33, 37, 38, 41, 46, 51, 74] learn image-to-image translation across a pair or several domains. UNIT [37] and MUNIT [20] learn shared representations between two domains for image translation. StarGAN [7, 8] learns a single network that takes in an input image and a style code to translate the input to multiple domains. SemanticGAN [34] jointly produces images and segmentation masks. On top of these, our model allows unique applications by exploiting geometry.

Leveraging Geometry. Keypoint representations have been used to find landmarks in an unsupervised way [44, 56, 57, 63, 69]. TransGaGa [63] uses a conditional VAE [32] to learn a heat map of facial landmarks to aid the image translation task. Jaderberg *et al.* [23] proposes a differentiable module to spatially modify feature maps. DeepWarp [13] learns to warp images for gaze manipulation. Caricature generation [6, 15, 55] has benefitted from warping the input photo for exaggerated facial features. However, they require supervision through paired caricature and photo data or facial feature detectors. In contrast, our generative model jointly models multiple domains by learning the geometric differences in a completely unsupervised manner.

3. Polymorphic-GAN

In Sec. 3.1 we describe the motivation for our proposed approach, the model architecture in Sec. 3.2 to 3.4 and the training procedure in Sec. 3.5.

3.1. Motivation

We aim to learn a GAN-based generator that can simultaneously synthesize *aligned* images from multiple different domains. We denote images as *aligned* if they share common attributes and conditions across domains, such as pose and lighting condition. To train such a model, it is critical that the features of the generator can be translated into each domain, and by sharing more layers in the generator, we can naturally enforce such alignment better. The intermediate features in modern GAN [5, 27, 28] generators are generally shaped as 3-dimensional tensors with spatial dimensions. They go through rendering layers such as shallow convolution layers to produce an output image.

Consider two image domains such as human faces and portrait paintings. They share many attributes, including geometry, and we can easily model both domains with shared generator layers and small domain-specific convolution layers that render facial features according to their domain. However, suppose the domains have a more significant gap in geometry, such as human and cat faces. In this case, the

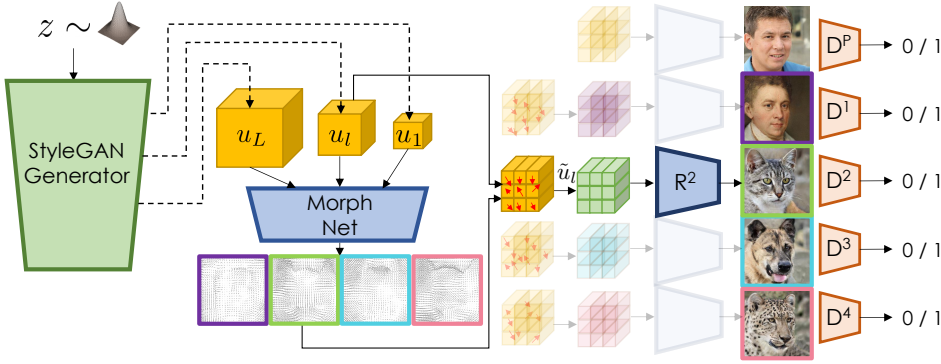


Figure 2. PMGAN generates aligned samples across multiple domains. All domains share the core StyleGAN generator. MorphNet produces domain-specific morph maps and warps the generator features to be geometrically suitable for the target domains. The learned morph maps can be exploited for interesting applications as we demonstrate in this work. We train the model with domain-specific discriminators.

rendering layers have to render geometrically different facial landmarks at different positions. The generator needs to learn the geometric differences by interpreting features in different ways for each domain, which discourages feature sharing. Therefore, we either need to reduce the number of shared layers in the generator or increase the number of domain-specific layers, either of which is not desirable as *alignment* between domains will be less enforced.

To overcome this problem, our PMGAN utilizes a domain-specific morph net that learns in a fully unsupervised manner the geometric differences between domains and morphs the shared features for each domain. It allows the sharing of entire generator layers across all domains while still having only shallow domain-specific rendering layers. We note that, in concurrent work, Wu *et al.* [65] analyzes how finetuning of a pre-trained StyleGAN2 model from a parent to a child domain affects the model’s network weights. The paper shows that the weights of the convolution layers in the main generator change the most, in addition to the mapping layers that produce the style vectors, which also change noticeably, especially when the geometric gap between the domains is large. These strong network parameter changes that are required to adapt the model indicate that naive feature sharing between domains, particularly geometrically different domains, is highly non-trivial. This is in line with our hypothesis. Our PMGAN more efficiently allows sharing of features across domains by explicitly modelling geometry in the generator’s feature space.

3.2. Pre-trained StyleGAN

PMGAN is based on StyleGAN2 [28] and we extend it to multiple domains (Figure 2). We denote the set of datasets for domains to be trained as $\mathcal{D} = \{\pi^P, \pi^1, \dots, \pi^N\}$ where π^P is a special dataset from the parent domain for which we assume that there exists a StyleGAN2 model pre-trained on π^P . PMGAN is composed of the pre-trained StyleGAN2’s generator G , domain-specific morph layers M^1, \dots, M^N and rendering layers R^1, \dots, R^N . We first sample a noise vector $z \sim p(z)$ from the standard Normal prior distribution and feed it through G , which produces the output image I^P and also the intermediate features u_1, \dots, u_L for L features in G .

Specifically, we store the generator features for each spatial resolution from $2^2 \times 2^2$ to $2^{L+1} \times 2^{L+1}$ before the final features are transformed via a 1×1 convolution layer (*i.e.* tRGB) that produces the output RGB values. We assume square images with $H = W = 2^{L+1}$.

3.3. MorphNet

Features u_1, \dots, u_L contain valuable information, including semantic content as well as fine-grained edge information. We use these features to produce domain-specific morph maps that can modify the geometry embedded in the features to be suitable for each target domain. The MorphNet component of PMGAN first reduces each feature map’s channel dimension to be smaller through a 1×1 convolution layer and then upsamples all features to match the largest spatial resolution $H \times W$. The upsampled features are concatenated channel-wise and go through two 3×3 convolution layers. We add a fixed 2-dimensional sinusoidal positional encoding [60] to the merged features to inject grid position information which can be useful for learning geometric biases in a dataset. Finally, this tensor is processed by domain-specific convolutional layers M^d for each domain d . M^d produces a $H \times W \times 2$ morph map \mathcal{M}_Δ^d , normalized between $[-1/\eta, 1/\eta]$ through Tanh activation function where η is a hyperparameter that controls the maximum displacement we allow the morphing operation to produce. \mathcal{M}_Δ^d represents the relative horizontal and vertical direction that each pixel would get its value from (a pixel here is (p, q) position in a 3-dim spatial tensor).

Algorithm 1 Inference step for PMGAN

```

function FORWARD( $z$ )
   $u_1, \dots, u_L, \mathcal{I}^P = G(z)$ 
   $u = \text{MergeFeatures}(u_1, \dots, u_L)$ 
  for  $d \in (1, \dots, N)$  do
     $\mathcal{M}_\Delta^d = M^d(u)$   $\triangleright$  Get morph map for domain  $d$ 
     $\{\tilde{u}_1, \dots, \tilde{u}_L\}_d = \text{Morph}(u_l, \mathcal{M}_\Delta^d)$  for all  $l$ 
     $\mathcal{I}^d = R^d(\tilde{u}_1, \dots, \tilde{u}_L)$ 
  return  $\mathcal{I}^1, \dots, \mathcal{I}^N, \mathcal{I}^P$ 

```

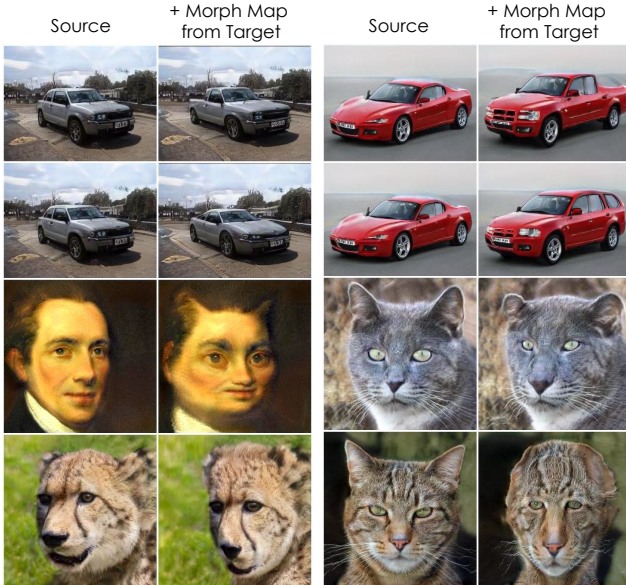


Figure 3. Effect of using target domain’s morph map while keeping everything else fixed in the source domain.

3.4. Feature Morphing

We follow Spatial Transformer Networks (SPN) [23] to differentially morph features with \mathcal{M}_Δ^d . We initialize a 2D sampling grid from an identity transformation matrix, normalized between $[-1, 1]$. The sampling grid has the same shape as \mathcal{M}_Δ^d , and each pixel (p, q) in the sampling grid contains the absolute position (x, y) of the source pixel that will be morphed into (p, q) . For example, if pixel (p, q) has value $(-1, -1)$, the vector at the top left corner of the source feature map will be morphed into (p, q) . The morph map \mathcal{M}_Δ^d is added to the grid, and we denote the resulting grid as $\Gamma \in \mathbb{R}^{H \times W \times 2}$. Unlike SPN that produces an affine transformation matrix with six parameters for sampling grid, we learn pixel-wise morphing maps, which gives us precise control for fine-detailed morphing. For each layer l of generator features, we perform the following Morph operation that bilinearly interpolates features:

$$\tilde{u}_l^{pq} = \sum_n^{H_l} \sum_m^{W_l} u_l^{nm} \max(0, 1 - |x^{pq} - m|) \max(0, 1 - |y^{pq} - n|) \quad (1)$$

where \tilde{u}_l^{pq} is the morphed feature vector at pixel (p, q) for layer l , u_l^{nm} is the source feature vector prior to Morph at pixel (n, m) , and (x^{pq}, y^{pq}) is the sample point in Γ for pixel (p, q) , assuming unnormalized grid coordinates for ease of presentation. Note that Γ is also bilinearly interpolated to match the spatial dimension of each layer (H_l, W_l) .

The morphed features $\{\tilde{u}_1, \dots, \tilde{u}_L\}_d$ are now geometrically transformed to be suitable for domain d . Each of these features is then processed via further convolution layers R^d to produce RGB images. They are finally summed together

using skip connections as in StyleGAN2 [28]. Importantly, the R^d layers can correct small unnatural distortions caused by the feature morphing process, in contrast to previous works that directly warp output images [55].

The inference step of PMGAN is summarized in Algo. 1.

3.5. Training

We use separate discriminators with the same architecture for each domain, and train PMGAN with non-saturating logistic loss [16], R1 regularization [42] and path-length regularization [28]. We use equal loss weightings for all domains, except when we do low-data regime training (Sec. 4.5), in which case we weigh losses by $|\pi^d|/\max_l |\pi^l|$ where $|\pi^d|$ is the number of training examples in domain d . The intuition is that we want the generator features to be mostly learned from data-rich domains while domains with significantly less data leverage the rich representation with domain-specific layers. The StyleGAN2 generator is initialized from pre-trained weights on a parent domain. We found that initializing all discriminators from the same pre-trained model helped stabilize training. We freeze the first three layers of discriminators and the shared generator, and do not update these weights [26, 43]. We also share the weights of k rendering layers of R across domains which promotes rendering of similar style such as colors. The more rendering layers we share, the more similar in style domains become, but that comes with the tradeoff of not being able to learn domain-specific styles. We found setting $k = 1$ or sharing the rendering layer at 4×4 spatial resolution was adequate in producing similarly styled outputs. We set the morph hyperparameter $\eta = 3$ for all experiments such that each pixel can move at most $1/6$ of the image size in the x and y direction. η can be adjusted depending on the geometric gap between domains.

4. Experiments

We construct two multi-domain datasets for evaluation:

Cars dataset consists of five classes of cars from the LSUN-Car dataset [67]. We use an object classifier by Ridnik *et al.* [50] that can output fine-grained object classes to divide the dataset into the following domains: Sedan (149K), SUV (52K), Sports car (58K), Van (25K), and Truck (22K), with the number of images in parentheses. The parent StyleGAN2 model is pre-trained on Sedan.

Faces dataset consists of Flickr-Faces-HQ [27] (70K), MetFaces [26] (1.3K), as well as Cat (5.6K), Dog (5.2K) and Wild life (5.2K) from the AFHQ dataset [8]. The parent model is pre-trained on the Flickr-Faces-HQ dataset.

AFHQ datasets have official testing splits, and we use 5% of the other domains as testing sets.

We carry out all experiments at 256×256 resolution. The datasets contain domains with varying geometric differences. Our goal is to learn both large and subtle geomet-

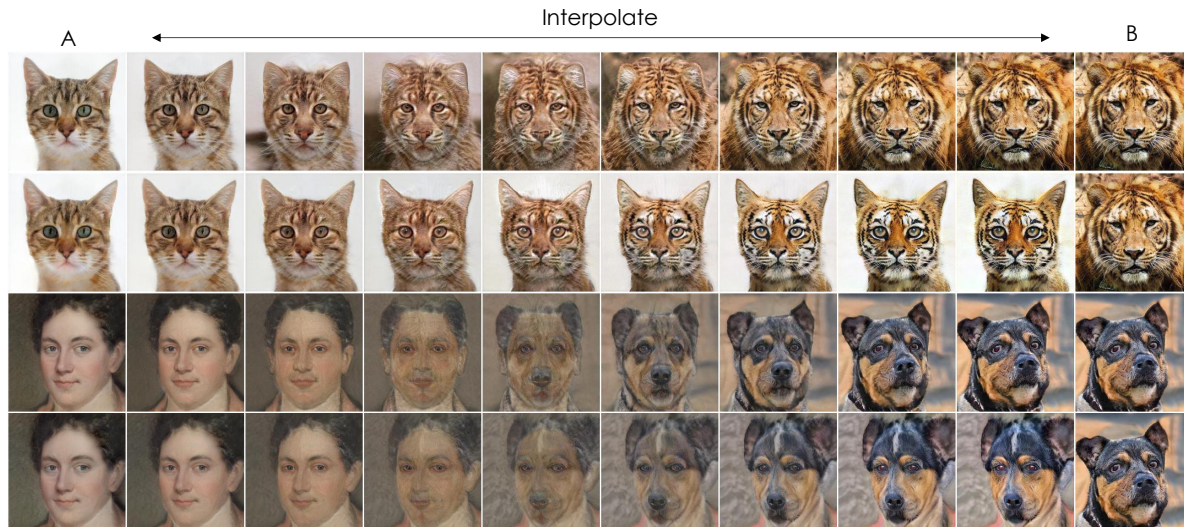


Figure 4. **Odd Rows:** Linearly interpolating both domain-specific layers of two domains and their latent vectors A&B. **Even Rows:** Linearly interpolating domain-specific layers and latent vectors while keeping the morph map of A the same.

Criterion	Method	Sedan	Truck	SUV	Sports Car	Van
FID (\downarrow)	*DC-StyleGAN2	18.0	120.1	189.1	80.1	111.3
	DC-StyleGAN2	10.3	93.4	82.6	93.3	96.5
	Ours w.o Morph	5.7	35.7	18.5	16.0	34.9
	Ours	5.4	23.3	11.3	9.1	19.3
Acc. (\uparrow)	*DC-StyleGAN2	85.9%	2.4%	6.3%	21.5%	16.2%
	DC-StyleGAN2	65.8%	3.8%	27.3%	14.0%	22.6%
	Ours w.o Morph	84.7%	45.2%	56.3%	75.3%	42.3%
	Ours	88.2%	69.2%	74.9%	88.8%	73.9%

Table 1. FID and classification accuracy for Cars dataset.

Criterion	Method	FFHQ	Metfaces	Cat	Dog	Wild Life
FID (\downarrow)	*DC-StyleGAN2	6.6	46.3	127.4	66.4	102.0
	DC-StyleGAN2	14.3	47.9	16.0	60.7	18.1
	Ours w.o Morph	8.1	37.5	21.3	63.8	27.8
	Ours	7.4	34.7	9.4	34.5	12.0
Acc. (\uparrow)	*DC-StyleGAN2	99.7%	87.8%	11.8%	77.0%	41.8%
	DC-StyleGAN2	92.5%	81.7%	96.2%	83.8%	87.3%
	Ours w.o Morph	99.9%	100.0%	96.6%	94.2%	98.8%
	Ours	99.9%	100.0%	99.5%	98.6%	99.7%

Table 2. FID and classification accuracy for Faces dataset.

rich and texture differences between them. We provide more details on each section in the supplementary materials.

4.1. Ablation Studies

We first verify the efficacy of PMGAN on producing aligned samples across domains with a single model. Our first baseline *Domain-Conditional StyleGAN2* (*DC-StyleGAN2*) is a modified StyleGAN2 model that takes a one-hot encoded domain vector as an input. The one-hot vector is embedded through a linear layer, concatenated with the output of the mapping network, merged with a linear layer and fed into the generator. **DC-StyleGAN2* has the same architecture as *DC-StyleGAN2*, but it starts from a pretrained model and only adds the extra layers for class-conditioning. The next one is *Ours without Morph*, the same as our full PMGAN, except for the MorphNet com-

ponent that morphs the generator features.

We measure sample quality with Fréchet Inception Distance (FID) [18] and accuracy using pretrained domain classifiers. The classifiers measure if models produce corresponding samples for each domain. They are implemented as ResNet-18 [17] for the 5-way classification task.

Tab. 1, 2 show ablation results. *DC-StyleGAN2* produces reasonable samples but they are not aligned well across domains as the domain conditioning in the generator modifies generator features to be specialized for each domain, as can be seen in Fig. 5. We found that **DC-StyleGAN2*, which finetunes a pre-trained model, has difficulties learning class-conditioning information, as can be seen in its low classification accuracy. We suspect that it is not trivial to adapt generator features to be suitable across domains without any domain-specific layers. *Ours without Morph* produces aligned poses as it tries to use the same generator features, but has trouble sharing features across domains because of their geometric differences. It cannot effectively use the shared features, as features corresponding to certain facial landmarks, such as eyes, nose and mouth, often vary in spatial position across domains.

In contrast, our full PMGAN leverages domain-specific layers but still benefits from sharing the entire stack of features due to the geometric morphing. It achieves the best overall sample quality and accuracy on both datasets.

4.2. Qualitative Analysis

We start by analyzing what PMGAN has learned in the morph maps. Similar to model interpolation [62, 65], we can interpolate the domain-specific layers of PMGAN to continuously interpolate two different domains. Additionally, we leverage the morph maps to investigate if the models correctly learned the geometric differences. We sample




Figure 5. Each row is a sample from one latent vector. *Top row*: *DC-StyleGAN2, *Second row*: DC-StyleGAN2, *Third row*: Ours without Morphing, *Last row*: Our PMGAN.

Criterion	Method	Truck	SUV	Sports Car	Van	Mean
mIoU (\uparrow)	Baseline	0.45	0.57	0.52	0.44	0.49
	Ours	0.67	0.74	0.63	0.64	0.67

Table 3. Mean IoU for zero-shot segmentation. Our transferred segmentation masks show high IoU with pseudo-labelled masks.

two latent vectors A and B, and linearly interpolate domain layer weights as well as the latent vectors. As can be seen in Fig. 4, PMGAN is capable of performing cross-domain interpolation, and by fixing A’s morph map during interpolation, it maintains the geometric characteristic of A while adapting to B’s texture. This shows how geometry is disentangled from rendering and PMGAN can be used for interesting image editing applications such as transforming a cat to look like a tiger. In Fig. 3, we show the effect of using the target domain t ’s morph map for a source domain s . Specifically, we swap the morph map \mathcal{M}_{Δ}^s with \mathcal{M}_{Δ}^t and render for domain s . For Cars, whose domains have similar texture, we can see how cars from source domains can be smoothly transformed towards the target domain. For Faces, we see interesting rendering such as a cat-shaped human face. These results demonstrate how PMGAN successfully learned the distinct geometries of each domain.

Edit Transfer There has been tremendous interest [36, 53, 54, 64] in disentangling StyleGAN’s latent space to find useful *edit vectors* that can modify the output image in a semantically meaningful way by pushing the latent vector of StyleGAN into certain directions. PMGAN’s aligned cross-domain samples through the shared generator allow us to discover edit vectors that transfer across domains. We use SeFa [54] for its simplicity to find edit vectors in PMGAN. We find meaningful vectors such as rotation, zoom, light-



Figure 6. Edit transfer. Edit directions discovered through PMGAN’s core generator can be transferred across all domains. Top: rotation, Middle: zoom, Bottom: color.




Figure 7. Zero-shot segmentation transfer. The masks in the left-most column are transferred to other domains using \mathcal{M}_{Δ} .

ing and elevation. Fig. 6 shows some examples of how edit vectors can be transferred across all domains.

4.3. Zero-shot Segmentation Transfer

Assuming there exists a method that can output a segmentation map for images from the parent domain, it is possible to zero-shot transfer the segmentation mask to all other domains using PMGAN’s learned morph map. We directly use the `Morph` operation on the segmentation map with




Figure 8. We compare image to image translation results from PMGAN and StarGANv2. PMGAN uses GAN inversion techniques to find a latent vector that can reconstruct the input image and renders the other domains.

\mathcal{M}_Δ after bilinearly interpolating the morph map to match the size of the mask. As the morph map \mathcal{M}_Δ captures the geometric differences between domains, we can successfully use \mathcal{M}_Δ to transfer the parent’s segmentation masks across domains as shown in Fig. 7. To measure the quality of the segmentation transfer, we use a pre-trained segmentation network to pseudo-label detailed car parts following Zhang *et al.* [70]. We compare the agreement between the pseudo-label and transferred segmentations from the sedan class. In Tab.3, the baseline measures mean IoU using the segmentation from the sedan class for all classes without morphing, which serves as a good baseline as PMGAN produces aligned samples whose poses are mostly identical. Our zero-shot segmentation shows much higher agreement with the pseudo-label, indicating our model correctly learned the correspondence between different car parts.

4.4. Image-to-Image Translation

There is a large body of work that does *GAN inversion* [1, 49, 58, 73] with StyleGAN. PMGAN can easily use any GAN inversion method as the model is based on StyleGAN. Once an image is inverted in the latent space, PMGAN can naturally be used for image-to-image translation (I2I) tasks by synthesizing every other domain with the same latent code. On the Cars dataset, we use latent optimization [27] to encode input images, and outperform the state-of-the-art multi-domain image translation model StarGANv2 [8] on both FID and accuracy. StarGANv2 has a strong shape bias from the input image and has trouble translating to another car domain, as indicated by its low accuracy in Tab. 4. For the Faces dataset, StarGANv2 does well if trained only on animal faces because geometric differences between animal classes are small. However, when trained on all five domains of Faces, training collapses and fails to translate between human and animal faces (Tab. 5 and Fig. 8). To compare with other image translation approaches, we also evaluate on a single domain translation task in Tab. 6. PMGAN shows competitive performance on the Cat-to-Dog task, despite being a generative model trained on all five domains together, as opposed to meth-

Criterion	Method	Sedan	Truck	SUV	Sports Car	Van	Mean
FID (↓)	StarGANv2 [8]	28.1	35.0	41.0	20.7	42.2	33.4
	Ours	26.7	25.3	26.6	25.6	35.9	28.0
Acc. (↑)	StarGANv2 [8]	48.5%	62.2%	58.1%	84.4%	63.5%	63.3%
	Ours	94.1%	90.0%	80.0%	91.6%	76.2%	86.4%

Table 4. I2I performance on Cars. Each column evaluates quality of samples translated from other domains to the column’s domain.

Evaluation Dataset	MUNIT [20]	DRIT [33]	MSGAN [41]	StarGANv2 [8]	Ours
Animals Only	41.5	95.6	61.4	16.2	33.1
All Domains	-	-	-	133.7	41.1

Table 5. I2I performance on Faces (FID). Ours is trained on all domains from Faces for both rows. Other models are trained only on animals for the first row, and on all domains for the second row.

Task	MUNIT [20]	CycleGAN [74]	StarGANv2 [8]	CUT [46]	Ours
Cat → Dog	91.4	76.3	53.4	56.4	55.9

Table 6. I2I performance on Cat-to-Dog (FID). Ours shows competitive performance despite being a generative model that jointly models all domains.

Dataset	Method	5%	20%	100%
Metfaces	StyleGAN2 [28] - single domain	68.7	83.0	72.4
	Ours - five domains	59.7	40.7	34.7
AFHQ-Cat	StyleGAN2 [28] - single domain	27.3	19.6	6.8
	Ours - five domains	23.3	13.8	9.4

Table 7. Low data regime (FID). For different amount of data used, we compare ours with StyleGAN2 trained on a single domain.

ods that only translate between two domains (except StarGANv2, which models three animal domains together).

4.5. Low Data Regime

PMGAN shares features for multiple domains, which can be beneficial for domains with small amounts of data, as they can leverage the rich representations learned from other domains. We evaluate PMGAN on the Faces dataset by varying the amount of data for the MetFaces and Cat domains while other domains use the full training data (Tab. 7). Compared to StyleGAN2, we achieve better FIDs when the amount of training data is small. Note that StyleGAN2 training with 5% data mode-collapsed. However, FID was not robust enough to reflect this, as the number of

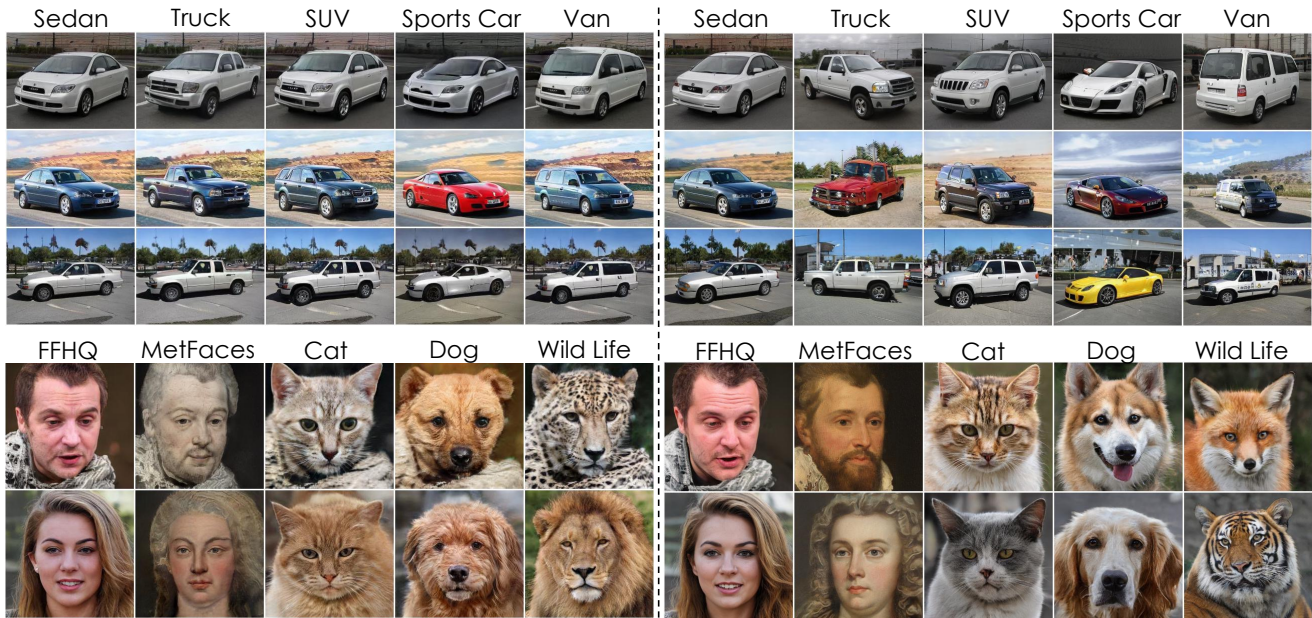


Figure 9. **Left:** Samples from PMGAN, **Right:** Samples from finetuned models from the same parent model. Our model produces consistently aligned samples across domains. Finetuning [43] specializes models for each domain, especially for less pre-aligned datasets.

data in MetFaces (1.3K) is too small. PMGAN can be combined with techniques that explicitly tackle low-data GAN training [26, 52, 71], which we leave for future work.

4.6. Comparison to Plain Fine-Tuning

There have been recent works [12, 45, 47, 65] on fine-tuning pre-trained StyleGANs for new target domains. While these methods can achieve high image quality, fine-tuning encourages the child models to be specialized to the new domains. As a further comparison, we fine-tune the same parent model used by our PMGAN for each domain [43]. For Faces, fine-tuning preserves some attributes such as pose and colors (with the same latents for original and fine-tuned models), but Fig. 9 shows PMGAN achieves better alignment in terms of facial shape and exact pose.

In contrast to the Faces data, Cars data has more diversity in viewpoints and car placement. The fine-tuned models show different sizes, poses and backgrounds. On the other hand, PMGAN produces consistently aligned cars. We evaluate the viewpoint alignment with the regression model from Liao *et al.* [35] by measuring the mean difference in azimuth and elevation between Sedan and other domains. Fine-tuning achieves 53.2 and 3.8 degrees in azimuth and elevation, respectively. PMGAN achieves **21.0** and **2.2** degrees in azimuth and elevation, significantly outperforming the fine-tuning approach. These results show that if domains have less diversity in poses and attributes, fine-tuning methods can produce reasonably aligned samples. However, as datasets become more diverse, it becomes challenging to enforce alignment without feature sharing. PMGAN has the unique advantage of being a model that

shares the *same features* across domains to produce highly aligned samples while enabling a diverse set of applications.

5. Limitations

We observed a slight deterioration in quality compared to StyleGAN2 on certain domains such as AFHQ-Cat (last column of Tab. 7). As we have used the same core generator backbone [28] for all experiments, future work includes improving the core generator, for example via increasing capacity to be more suitable for multi-domain modelling. For *Cars*, PMGAN often puts vibrant colors on *Sports Car* as it consists of mostly those colors. One possible remedy is adding a regularization term encouraging each domain to output similar colors. Lastly, the morph map in PMGAN is 2D and consequently cannot handle morphing in 3D. If the geometric differences between domains have to be modelled in 3D, such as object rotations, PMGAN can only mimic them rather than performing true 3D morphing. PMGAN is also not applicable for vastly different domains.

6. Conclusion

We introduced *Polymorphic-GAN*, which produces aligned samples across multiple domains by learning the geometric differences through morph maps. PMGAN’s morph maps enable efficient sharing of generator features. This allows PMGAN to be utilized for diverse applications, including zero-shot segmentation transfer, image-to-image translation, image editing across multiple domains as well as training in low-data settings. PMGAN is the first GAN to efficiently synthesize aligned samples from multiple geometrically-varying domains at the same time.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2stylegan: How to embed images into the stylegan latent space? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4432–4441, 2019. 1, 2, 7
- [2] Kyungjune Baek, Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Hyunjung Shim. Rethinking the truly unsupervised image-to-image translation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14154–14163, 2021. 2
- [3] David Bau, Hendrik Strobelt, William Peebles, Jonas Wulff, Bolei Zhou, Jun-Yan Zhu, and Antonio Torralba. Semantic photo manipulation with a generative image prior. *ACM Trans. Graph.*, 38(4), 2019. 1
- [4] David Bau, Jun-Yan Zhu, Hendrik Strobelt, Bolei Zhou, Joshua B. Tenenbaum, William T. Freeman, and Antonio Torralba. Gan dissection: Visualizing and understanding generative adversarial networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019. 1
- [5] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 1, 2
- [6] Kaidi Cao, Jing Liao, and Lu Yuan. Carigans: Unpaired photo-to-caricature translation. *ACM Transactions on Graphics (Proc. of Siggraph Asia 2018)*, 2018. 2
- [7] Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 1, 2
- [8] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 1, 2, 4, 7
- [9] Edo Collins, Raja Bala, Bob Price, and Sabine Süsstrunk. Editing in style: Uncovering the local semantics of GANs. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [10] Jifeng Dai, Kaiming He, and Jian Sun. Instance-aware semantic segmentation via multi-task network cascades. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3150–3158, 2016. 2
- [11] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12873–12883, 2021. 1
- [12] Rinon Gal, Or Patashnik, Haggai Maron, Gal Chechik, and Daniel Cohen-Or. Stylegan-nada: Clip-guided domain adaptation of image generators. *arXiv preprint arXiv:2108.00946*, 2021. 2, 8
- [13] Yaroslav Ganin, Daniil Kononenko, Diana Sungatullina, and Victor Lempitsky. Deepwarp: Photorealistic image resynthesis for gaze manipulation. In *European conference on computer vision*, pages 311–326. Springer, 2016. 2
- [14] Lore Goetschalckx, Alex Andonian, Aude Oliva, and Phillip Isola. Ganalyze: Toward visual definitions of cognitive image properties. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 1
- [15] Julia Gong, Yannick Hold-Geoffroy, and Jingwan Lu. Auto-ton: Automatic geometric warping for face cartoon generation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 360–369, 2020. 2
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 4
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [18] 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. *Advances in neural information processing systems*, 30, 2017. 5
- [19] Xianxu Hou, Xiaokang Zhang, Linlin Shen, Zhihui Lai, and Jun Wan. Guidedstyle: Attribute knowledge guided style manipulation for semantic face editing. *arXiv preprint arXiv:2012.11856*, 2020. 1
- [20] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 172–189, 2018. 2, 7
- [21] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. In *Proc. NeurIPS*, 2020. 1
- [22] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. 1
- [23] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. *NeurIPS*, 28:2017–2025, 2015. 2, 4
- [24] Ali Jahanian*, Lucy Chai*, and Phillip Isola. On the ”steerability” of generative adversarial networks. In *International Conference on Learning Representations*, 2020. 1
- [25] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *ICLR*, 2018. 1
- [26] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Proc. NeurIPS*, 2020. 1, 2, 4, 8
- [27] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. 1, 2, 4, 7

- [28] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 1, 2, 3, 4, 7, 8
- [29] H. Kazemi, S. Iranmanesh, and N. Nasrabadi. Style and content disentanglement in generative adversarial networks. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 848–856, Los Alamitos, CA, USA, Jan 2019. IEEE Computer Society. 1
- [30] Hyunsu Kim, Yunjey Choi, Junho Kim, Sungjoo Yoo, and Youngjung Uh. Exploiting spatial dimensions of latent in gan for real-time image editing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 1
- [31] Seung Wook Kim, Jonah Philion, Antonio Torralba, and Sanja Fidler. DriveGAN: Towards a Controllable High-Quality Neural Simulation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1
- [32] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2
- [33] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse image-to-image translation via disentangled representations. In *Proceedings of the European conference on computer vision (ECCV)*, pages 35–51, 2018. 2, 7
- [34] Daiqing Li, Junlin Yang, Karsten Kreis, Antonio Torralba, and Sanja Fidler. Semantic segmentation with generative models: Semi-supervised learning and strong out-of-domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8300–8311, 2021. 1, 2
- [35] Shuai Liao, Efstratios Gavves, and Cees G. M. Snoek. Spherical regression: Learning viewpoints, surface normals and 3d rotations on n-spheres. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Long Beach, USA, June 2019. 8
- [36] Huan Ling, Karsten Kreis, Daiqing Li, Seung Wook Kim, Antonio Torralba, and Sanja Fidler. Editgan: High-precision semantic image editing. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 1, 2, 6
- [37] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *Advances in neural information processing systems*, pages 700–708, 2017. 2
- [38] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. Few-shot unsupervised image-to-image translation. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 2
- [39] Shengchao Liu, Yingyu Liang, and Anthony Gitter. Loss-balanced task weighting to reduce negative transfer in multi-task learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9977–9978, 2019. 2
- [40] Jiaqi Ma, Zhe Zhao, Xinyang Yi, Jilin Chen, Lichan Hong, and Ed H Chi. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1930–1939, 2018. 2
- [41] Qi Mao, Hsin-Ying Lee, Hung-Yu Tseng, Siwei Ma, and Ming-Hsuan Yang. Mode seeking generative adversarial networks for diverse image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1429–1437, 2019. 2, 7
- [42] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *International conference on machine learning*, pages 3481–3490. PMLR, 2018. 4
- [43] Sangwoo Mo, Minsu Cho, and Jinwoo Shin. Freeze the discriminator: a simple baseline for fine-tuning gans. In *CVPR AI for Content Creation Workshop*, 2020. 1, 2, 4, 8
- [44] Ron Mokady, Rotem Tzaban, Sagie Benaim, Amit H Bermano, and Daniel Cohen-Or. Jokr: Joint keypoint representation for unsupervised cross-domain motion retargeting. *arXiv preprint arXiv:2106.09679*, 2021. 2
- [45] Utkarsh Ojha, Yijun Li, Jingwan Lu, Alexei A Efros, Yong Jae Lee, Eli Shechtman, and Richard Zhang. Few-shot image generation via cross-domain correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10743–10752, 2021. 2, 8
- [46] Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *European Conference on Computer Vision*, pages 319–345. Springer, 2020. 2, 7
- [47] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2085–2094, 2021. 2, 8
- [48] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021. 2
- [49] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2287–2296, 2021. 1, 2, 7
- [50] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihl Zelnik-Manor. Imagenet-21k pretraining for the masses. *arXiv preprint arXiv:2104.10972*, 2021. 4
- [51] Kuniaki Saito, Kate Saenko, and Ming-Yu Liu. Cocomunit: Few-shot unsupervised image translation with a content conditioned style encoder. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 382–398. Springer, 2020. 2
- [52] Axel Sauer, Kashyap Chitta, Jens Müller, and Andreas Geiger. Projected gans converge faster. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 8
- [53] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou. Interfacegan: Interpreting the disentangled face representa-

- tion learned by gans. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 2, 6
- [54] Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1532–1540, 2021. 1, 2, 6
- [55] Yichun Shi, Debayan Deb, and Anil K Jain. Warpgan: Automatic caricature generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10762–10771, 2019. 2, 4
- [56] Supasorn Suwajanakorn, Noah Snively, Jonathan Tompson, and Mohammad Norouzi. Discovery of latent 3d keypoints via end-to-end geometric reasoning. *arXiv preprint arXiv:1807.03146*, 2018. 2
- [57] James Thewlis, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object landmarks by factorized spatial embeddings. In *Proceedings of the IEEE international conference on computer vision*, pages 5916–5925, 2017. 2
- [58] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021. 1, 2, 7
- [59] Nontawat Tritrong, Pitchaporn Rewatbowornwong, and Supasorn Suwajanakorn. Repurposing gans for one-shot semantic part segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4475–4485, 2021. 1
- [60] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 3
- [61] Andrey Voynov and Artem Babenko. Unsupervised discovery of interpretable directions in the gan latent space. In *International Conference on Machine Learning*, pages 9786–9796. PMLR, 2020. 1
- [62] Xintao Wang, Ke Yu, Chao Dong, Xiaoou Tang, and Chen Change Loy. Deep network interpolation for continuous imagery effect transition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1692–1701, 2019. 5
- [63] Wayne Wu, Kaidi Cao, Cheng Li, Chen Qian, and Chen Change Loy. Transgaga: Geometry-aware unsupervised image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8012–8021, 2019. 2
- [64] Zongze Wu, Dani Lischinski, and Eli Shechtman. Stylespace analysis: Disentangled controls for stylegan image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12863–12872, 2021. 1, 2, 6
- [65] Zongze Wu, Yotam Nitzan, Eli Shechtman, and Dani Lischinski. Stylealign: Analysis and applications of aligned stylegan models. *arXiv preprint arXiv:2110.11323*, 2021. 3, 5, 8
- [66] Jianjin Xu and Changxi Zheng. Linear semantics in generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9351–9360, 2021. 1
- [67] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 4
- [68] Yuxuan Zhang, Wenzheng Chen, Huan Ling, Jun Gao, Yanan Zhang, Antonio Torralba, and Sanja Fidler. Image gans meet differentiable rendering for inverse graphics and interpretable 3d neural rendering. *arXiv preprint arXiv:2010.09125*, 2020. 1
- [69] Yuting Zhang, Yijie Guo, Yixin Jin, Yijun Luo, Zhiyuan He, and Honglak Lee. Unsupervised discovery of object landmarks as structural representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2694–2703, 2018. 2
- [70] Yuxuan Zhang, Huan Ling, Jun Gao, Kangxue Yin, Jean-Francois Lafleche, Adela Barriuso, Antonio Torralba, and Sanja Fidler. Datasetgan: Efficient labeled data factory with minimal human effort. In *CVPR*, 2021. 1, 2, 7
- [71] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient gan training. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020. 8
- [72] Xiangyun Zhao, Haoxiang Li, Xiaohui Shen, Xiaodan Liang, and Ying Wu. A modulation module for multi-task learning with applications in image retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 401–416, 2018. 2
- [73] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *European conference on computer vision*, pages 592–608. Springer, 2020. 1, 2, 7
- [74] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 2, 7