

# Data InStance Prior (DISP) in Generative Adversarial Networks

Puneet Mangla<sup>\*3 †</sup>      Nupur Kumari<sup>\*2 ‡</sup>      Mayank Singh<sup>\*2 †</sup>  
 pmangla261@gmail.com      nupurkmr9@gmail.com      mayanksingh027@gmail.com

Balaji Krishnamurthy<sup>1</sup>      Vineeth N Balasubramanian<sup>3</sup>  
 kbalaji@adobe.com      vineethnb@iith.ac.in

<sup>1</sup>Media and Data Science Research lab, Adobe    <sup>2</sup>CMU    <sup>3</sup>IIT Hyderabad, India

## Abstract

*Recent advances in generative adversarial networks (GANs) have shown remarkable progress in generating high-quality images. However, this gain in performance depends on the availability of a large amount of training data. In limited data regimes, training typically diverges, and therefore the generated samples are of low quality and lack diversity. Previous works have addressed training in low data setting by leveraging transfer learning and data augmentation techniques. We propose a novel transfer learning method for GANs in the limited data domain by leveraging informative data prior derived from self-supervised/supervised pre-trained networks trained on a diverse source domain. We perform experiments on several standard vision datasets using various GAN architectures (BigGAN, SNGAN, StyleGAN2) to demonstrate that the proposed method effectively transfers knowledge to domains with few target images, outperforming existing state-of-the-art techniques in terms of image quality and diversity. We also show the utility of data instance prior in large-scale unconditional image generation.*

## 1. Introduction

Generative Adversarial Networks (GANs) are at the forefront of modern high-quality image synthesis in recent years [4, 24, 23]. GANs have also demonstrated excellent performance on many related computer vision tasks such as image manipulation [70, 17], image editing [41, 48, 18], and compression [51]. Despite the success in large-scale image synthesis, GAN training suffers from a several drawbacks that arise in practice, such as training instability and mode

collapse [11, 2]. This has prompted research in several non-adversarial generative models [16, 3, 28, 25]. These techniques are implicitly designed to overcome the mode collapse problem, however, the quality of generated samples is still not on par with GANs.

Current state-of-the-art deep generative models require a large volume of data and computation resources. The collection of large datasets of images suitable for training - especially labeled data in case of conditional GANs - can also be a daunting task due to issues such as copyright, image quality. To curb these limitations, researchers have recently proposed techniques inspired by transfer learning [37, 57, 36] and data augmentation [21, 66, 63]. Advancements in data and computation efficiency for image synthesis can enable its applications in data-deficient fields such as medicine [59] where labeled data procurement can be difficult.

Transfer learning is a promising area of research [39, 40] that leverages prior information acquired from large datasets to help in training models on a target dataset under limited data and resource constraints. There has been extensive exploration of transfer learning in classification problems that have shown excellent performance on various downstream data-deficient domains. Similar extensions of reusing pre-trained networks for transfer learning (i.e. fine-tuning a subset of pre-trained network weights from a data-rich domain) have also been recently employed for image synthesis in GANs [57, 37, 36, 56, 65] in the limited data regime. However, these approaches are still prone to overfitting on the sparse target data, and hence suffer from degraded image quality and diversity.

In this work, we propose a simple yet effective way of transferring prior knowledge in unsupervised image generation given a small sample size ( $\sim 100$ -2000) of the target data distribution. Our approach is motivated by the formulation of the IMLE technique [28] that seeks to obtain mode coverage of target data distribution by learning a mapping

\* Authors contributed equally

<sup>†</sup> Work done during Adobe MDSR internship

<sup>‡</sup> Work done while at Adobe

between latent and target distributions using a maximum likelihood criterion. We instead propose the use of data priors in GANs to match the representation of the generated samples to real modes of data. In particular, we show that using an informative *data instance prior* in limited and large-scale unsupervised image generation substantially improves the performance of image synthesis. We show that these data priors can be derived from commonly used computer vision pre-trained networks [49, 64, 37, 16] or self-supervised data representations [5] (without any violation of the target setting’s requirements, i.e. ensuring that the pre-trained network has not been trained on few-shot classes in the few-shot learning setting, for instance). In case of sparse training data, our approach of using data instance priors leverages a model pre-trained on a rich source domain to learn the target distribution. Different from previous works [37, 56, 57] which rely on fine-tuning models trained on a data-rich domain, we propose to leverage the feature representations of the source model as data instance priors, to distill knowledge [44, 15] into the target generative problem setting.

We note that our technique of using instance level priors for transfer learning becomes fully unsupervised in case the data priors are extracted from self-supervised pre-trained networks. Furthermore, in addition to image generation in low data domain, we also achieve state-of-the-art Fréchet inception distance (FID) score [14] on large-scale unsupervised image generation.

We summarize our main contributions as follows:

- We propose Data InStance Prior (DISP), a novel transfer learning technique for GAN image synthesis in low-data regime. We show that employing DISP in conjunction with existing few-shot image generation methods outperforms state-of-the-art results. We show with as little as 100 images our approach DISP results in generation of diverse and high quality images (see Figure 3).
- We demonstrate the utility of our approach in large-scale unsupervised GANs [34, 4]. It achieves the new state-of-the-art in terms of image quality [14] and diversity [46, 33].

We call our method *data instance prior* (and not just data prior), since it uses representations of instances as a prior, and not a data distribution itself.

## 2. Related Work

**Deep Generative Models** In recent years, there has been a surge in the research of deep generative models. Some of the popular approaches include variational auto-encoders (VAEs) [43, 25], auto-regressive (AR) models [54, 53] and GANs [12]. VAE models learn by maximizing the variational lower bound of training data likelihood. Auto-regressive approaches model the data distribution as a prod-

uct of the conditional probabilities to sequentially generate data. GANs comprise of two networks, a generator and a discriminator that train in a min-max optimization. Specifically, the generator aims to generate samples to fool the discriminator, while the discriminator learns to distinguish these generated samples from the real samples. Several research efforts in GANs have focused on improving the performance [20, 8, 42, 24, 23, 4, 63] and training stability [47, 13, 1, 34, 31, 6]. Recently, the areas of latent space manipulation for semantic editing [48, 18, 69, 41] and few-shot image generation [56, 36, 37] have gained traction in an effort to mitigate the practical challenges while deploying GANs. Several other non-adversarial training approaches such as [16, 3, 28, 62] have also been explored for generative modeling, which leverage supervised learning along with perceptual loss [64] for training such models.

**Transfer Learning in GANs** While there has been extensive research in the area of transfer learning for classification models [60, 39, 52, 40, 9], relatively fewer efforts have explored this on the task of data generation [57, 56, 37, 65, 36]. [57] proposed to fine-tune a pre-trained GAN model (often having millions of parameters) from a data-rich source to adapt to the target domain with limited samples. This approach, however, often suffers from overfitting as the final model parameters are updated using only few samples of the target domain. To counter overfitting, the work of [37] proposes to update only the batch normalization parameters of the pre-trained GAN model. In this approach, however, the generator is not adversarially trained and uses supervised  $L_1$  pixel distance loss and perceptual loss [19, 64] which often leads to generation of blurry images in the target domain. Based on the assumption that source and target domain support sets are similar, [56] recently proposed to learn an additional mapping network that transforms the latent code suitable for generating images of target domain while keeping the other parameters frozen. We show that our method DISP outperforms the leading baselines in few-shot image generation including [37, 56, 66].

A related line of recent research aims to improve large-scale unsupervised image generation in GANs by employing self-supervision - in particular, an auxiliary task of rotation prediction [6] or using one-hot labels obtained by clustering in the discriminator’s [29] or ImageNet classifier feature space [45]. In contrast, our method utilizes instance level priors derived from the feature activations of self-supervised/supervised pre-trained networks to improve unconditional few-shot and large-scale image generation, leading to simpler formulation and better performance as shown in our experiments in Section 5.3 and Supplementary. Recently, some methods [21, 66, 63, 67] have leveraged data augmentation to effectively increase the number of samples and prevent overfitting in GAN training. How-

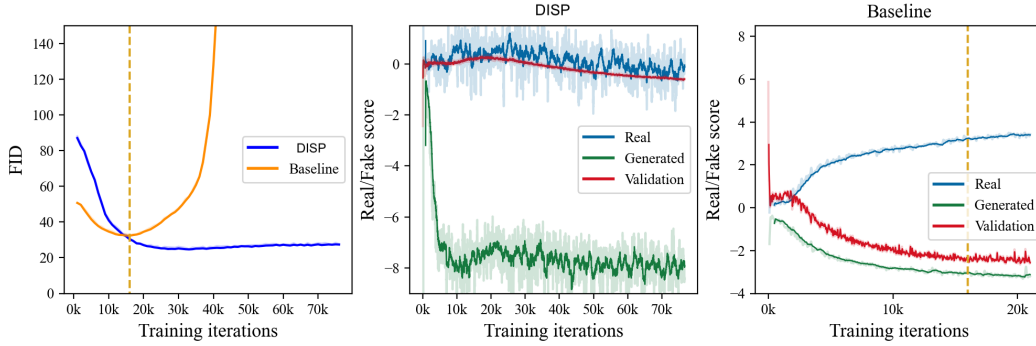


Figure 1: **Discriminator overfitting in limited data.** *left*: when trained on 10% data of CIFAR-100 FID of baseline model on training set increases very early in training(at around 15k iterations) unlike the FID curve of DISP. *middle*: Discriminator score on training and validation images remain similar to each other and consistently higher than score of generated images for DISP model. *right*: Discriminator score on training and validation images diverges and the training collapses for the baseline model.

ever, data augmentation techniques often times alter the true data distribution and there is a leakage of these augmentations to the generated image, as shown in [67, 66]. To overcome this, [66] recently proposed to use differential augmentation and [21] leveraged an adaptive discriminator augmentation mechanism. We instead focus on leveraging informative instance level priors and also show how our method can be used in conjunction with augmentation techniques [66] to further improve the performance.

### 3. Preliminaries

We briefly describe Conditional Generative Networks(cGANs) before discussing our methodology. cGANs consists of a generator network  $G$  which is trained adversarially with a discriminator network  $D$  to learn a target data distribution  $q(\mathbf{x}|y)$ . Given a noise vector  $z$  and a condition vector  $y$ ,  $G$  generates a sample  $x$  e.g. an image and the role of  $D$  is to distinguish between real samples and those generated from  $G$ . Conditional GANs use auxiliary information  $y$  for e.g. class label of the sample as input in the generator and discriminator networks. The standard hinge loss [50] for training cGANs is given by:

$$\begin{aligned}
 L_D &= \mathbb{E}_{y \sim q(y)} \left[ \mathbb{E}_{\mathbf{x} \sim q(\mathbf{x}|y)} [\max(0, 1 - D(\mathbf{x}, y))] \right] \\
 &\quad + \mathbb{E}_{y \sim q(y)} \left[ \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} [\max(0, 1 + D(G(\mathbf{z}|y), y))] \right] \quad (1) \\
 L_G &= -\mathbb{E}_{y \sim q(y)} \left[ \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} [D(G(\mathbf{z}|y), y)] \right]
 \end{aligned}$$

where the discriminator score  $D(\mathbf{x}, y)$  depends on input image (either real or fake) and conditional label  $y$  [35, 38]. The label information is generally passed into  $G$  through a one-hot vector concatenated with  $z$  or through conditional batch norm layers [7, 10].

## 4. Methodology

We propose a transfer learning framework, Data Instance Prior (DISP), for training GANs that exploits knowledge extracted from self-supervised/supervised networks, pre-trained on a rich and diverse source domain in the form of instance level priors. GANs are observed to be prone to mode-collapse that is further exacerbated in case of sparse training data. It has been shown that providing class label information in GANs significantly improves training stability and quality of generated images as compared to unconditional setting [35, 6]. We take motivation from the reconstructive framework of IMLE [28] and propose to condition GANs on image instance prior that acts as a regularizer to prevent mode collapse and discriminator overfitting.

**Knowledge Transfer in GAN** GANs are a class of implicit generative models that minimize a divergence measure between the data distribution  $q(\mathbf{x})$  and the generator output distribution  $G(\mathbf{z})$  where  $\mathbf{z} \sim p(\mathbf{z})$  denotes the latent distribution. Intuitively, the minimization of this divergence objective ensures that each generated sample  $G(\mathbf{z})$  is close to some data example  $\mathbf{x} \sim q(\mathbf{x})$ . However, this does not ensure the converse, i.e. each real sample has a generated sample in its vicinity, which can result in mode dropping. To counter this, especially in limited data regime, we propose to update the parameters of the model so that each real data example is close to some generated sample similar to [28] by using data instance priors as conditional label in GANs. We use the image features extracted from networks pre-trained on a large source domain as the instance level prior to enable transfer of knowledge.

Given a pre-trained feature extractor  $C : \mathbb{R}^p \rightarrow \mathbb{R}^d$ , which is trained on a source domain using supervisory signals or self-supervision, we use its output  $C(\mathbf{x})$  as the conditional information during GAN training.  $G$  is conditioned

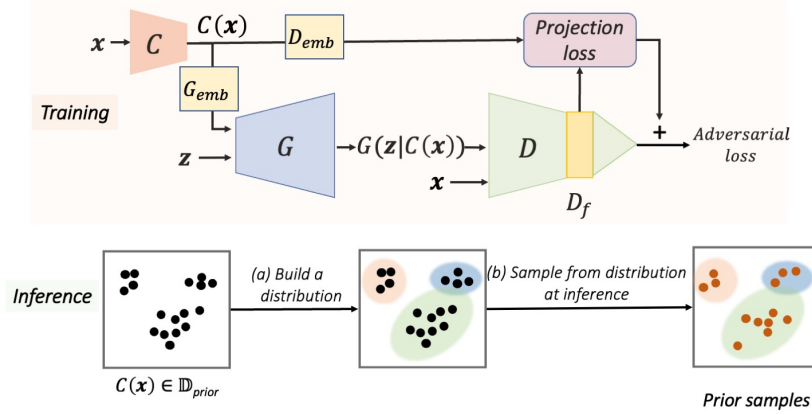


Figure 2: Overview of our proposed technique, Data Instance Priors (DISP) for transfer learning in GANs. *Top:* DISP training with feature  $C(\mathbf{x})$  of a real sample  $\mathbf{x}$  as a conditional prior in the conditional GAN framework of [35].  $C$  is a pre-trained network on a rich source domain from which we wish to transfer knowledge. *Bottom:* Inference over trained GAN involves learning a distribution over the set of training data prior  $\{C(\mathbf{x})\}$  to enable sampling of conditional priors.

on  $C(\mathbf{x})$  using conditional batch-norm [10] whose input is  $G_{emb}(C(\mathbf{x}))$ , where  $G_{emb}$  is a learnable projection matrix. During training we enforce that  $G(\mathbf{z}|C(\mathbf{x}))$  is close to the real image  $\mathbf{x}$  in discriminator feature space. Let the discriminator be  $D = D_l \circ D_f$  ( $\circ$  denotes composition) where  $D_f$  is discriminator’s last feature layer and  $D_l$  is the final linear classifier layer. To enforce the above objective we map  $C(\mathbf{x})$  to discriminator’s feature layer dimension using a trainable projection matrix  $D_{emb}$  and minimize distance between  $D_{emb}(C(\mathbf{x}))$  and  $D_f$  of both real image  $\mathbf{x}$  and generated image  $G(\mathbf{z}|C(\mathbf{x}))$  in an adversarial manner. Hence, our final GAN training loss for the discriminator and generator is given by:

$$\begin{aligned}
 L_D &= \mathbb{E}_{\mathbf{x} \sim q(x)} [\max(0, 1 - D(\mathbf{x}, C(\mathbf{x}))) \\
 &\quad + \mathbb{E}_{\mathbf{x} \sim q(x), \mathbf{z} \sim p(\mathbf{z})} [\max(0, 1 + D(G(\mathbf{z}|C(\mathbf{x})), C(\mathbf{x}))) \\
 L_G &= - \mathbb{E}_{\mathbf{x} \sim q(x), \mathbf{z} \sim p(\mathbf{z})} [D(G(\mathbf{z}|C(\mathbf{x})), C(\mathbf{x}))]
 \end{aligned} \tag{2}$$

where

$$D(\mathbf{x}, \mathbf{y}) = D_{emb}(\mathbf{y}) \cdot D_f(\mathbf{x}) + D_l \circ D_f(\mathbf{x}) \tag{3}$$

In the above formulation, the first term in Eq. 3 is the projection loss as in [35] between input image and conditional embedding of discriminator. Since conditional embedding is extracted from a pre-trained network, above training objective leads to feature level knowledge distillation from  $C$ . It also acts as a regularizer on the discriminator reducing its overfitting in the limited data setting. As shown in Figure 1, the gap between discriminator score ( $D_l \circ D_f$ ) of training and validation images keeps on increasing and FID quickly degrades for baseline model as compared to DISP when

trained on only 10% data of CIFAR-100. Moreover, enforcing feature  $D_f(G(\mathbf{z}|C(\mathbf{x})))$  to be similar to  $D_{emb}(C(\mathbf{x}))$  promotes that for each real sample, there exists a generated sample close to it and hence promotes mode coverage of target data distribution. We demonstrate that the above proposed use of data instance priors from a pre-trained feature extractor, while designed for a limited data setting, also benefits in large-scale image generation. Our overall methodology is illustrated in Figure 2 and pseudo code is provided in the Supplementary section.

**Random image generation at inference** Given the training set  $\mathbb{D}_{image} = \{\mathbf{x}_j\}_{j=1}^n$  of sample size  $n$  and its corresponding training data prior set  $\mathbb{D}_{prior} = \{C(\mathbf{x}_j)\}_{j=1}^n$ , the generator requires access to  $\mathbb{D}_{prior}$  for sample generation. In case of few-shot and limited image generation where size of  $\mathbb{D}_{prior}$  is less, to create more variations, we generate images conditioned on prior samples from a vicinal mix distribution i.e

$$G(\mathbf{z}|\mathbf{p}) \quad \text{where } \mathbf{p} \sim \mathcal{V}_{mix} \tag{4}$$

The vicinal mix distribution is defined as:

$$\mathcal{V}_{mix}(\mathbf{p}) = \frac{1}{|\mathbb{D}_{prior}|^2} \sum_{i,j} \mathbb{E}_\lambda [\delta(\lambda \cdot \mathbf{p}_i + (1 - \lambda) \cdot \mathbf{p}_j)]$$

where  $\lambda \sim \mathcal{U}[0, 1]$  and  $\delta(\cdot)$  is dirac-delta function

(5)

In case of large-scale image generation, we learn a Gaussian Mixture Model (GMM) [58] on  $\mathbb{D}_{prior}$ . This enables memory efficient sampling of conditional prior from the learned GMM distribution during inference:

$$G(\mathbf{z}|\mathcal{N}(\mu, \Sigma)) \quad \text{where } \mu, \Sigma \sim \text{GMM}(G_{emb}(\mathbb{D}_{prior})) \tag{6}$$

Method	Pre-training	SNGAN (128 x 128)					
		Anime			Faces		
		FID ↓	P ↑	R ↑	FID ↓	P ↑	R ↑
From scratch + DISP-Vgg16	✗	120.38 <b>66.85</b>	0.61 <b>0.71</b>	0.00 <b>0.03</b>	140.66 <b>68.49</b>	0.31 <b>0.74</b>	0.00 <b>0.15</b>
TransferGAN + DISP-Vgg16	✓	102.75 <b>86.96</b>	<b>0.70</b> 0.57	0.00 <b>0.02</b>	101.15 <b>75.21</b>	<b>0.85</b> 0.70	0.00 <b>0.10</b>
FreezeD + DISP-Vgg16 + DISP-SimCLR	✓	109.40 93.36 <b>89.39</b>	<b>0.67</b> 0.56 0.46	0.00 <b>0.03</b> 0.025	107.83 77.09 <b>70.40</b>	<b>0.83</b> 0.68 0.74	0.00 0.14 <b>0.22</b>
ADA + DISP-Vgg16	✗	78.28 <b>60.8</b>	0.87 <b>0.90</b>	0.0 <b>0.003</b>	159.3 <b>79.5</b>	0.69 <b>0.85</b>	0.0 <b>0.004</b>
DiffAugment + DISP-Vgg16 + DISP-SimCLR	✗	85.16 <b>48.67</b> 52.41	<b>0.95</b> 0.82 0.77	0.00 0.03 <b>0.04</b>	109.25 <b>62.44</b> 64.53	<b>0.84</b> 0.80 0.78	0.00 0.19 <b>0.22</b>
BSA* GLANN + DISP-Vgg16	✓	92.0 <b>67.07</b>	- <b>0.87</b>	- <b>0.01</b>	123.2 <b>60.11</b>	- <b>0.95</b>	- <b>0.08</b>

Table 1: Few-shot image generation performance using 100 training images (↓: lower is better; ↑: higher is better). Precision and Recall scores are based on [27]. FID is computed using 10k, 7k generated and 10k, 7k real samples for Anime and Faces respectively. \* denotes directly reported from the paper.

## 5. Experiments

We perform extensive experiments to highlight the efficacy of our data instance prior module, DISP in unsupervised training based on SNGAN [34], BigGAN [4] and StyleGAN2 [24] architectures. For extracting image prior information, we use the last layer activations of: Vgg16 [49] classification network trained on ImageNet; and SimCLR [5] network trained using self-supervision on ImageNet. We conduct experiments on (1) *few-shot* ( $\sim 25$ -100 images), (2) *limited* ( $\sim 2k$ -5k images) and (3) *large-scale* ( $\sim 50k$ -1M images) data settings. For evaluation, we use FID [14], precision and recall scores [27] to assess the quality and mode-coverage/diversity of the generated images.

### 5.1. Few-Shot Image Generation

**Baselines and Datasets** We compare and augment our methodology DISP with training SNGAN from scratch and the following leading baselines: Batch Statistics Adaptation (BSA) [37], TransferGAN [57], FreezeD [36], ADA [22] and DiffAugment [66]. In case of BSA, a non-adversarial variant, GLANN [16] is used which optimizes for image embeddings and generative model through perceptual loss<sup>1</sup>. We use our data priors to distill knowledge over these image embeddings. For more training and hyperparameter details, please refer to Supplementary.

We perform experiments on randomly chosen 100 images at  $128 \times 128$  resolution from: (1) Anime<sup>2</sup> and (2)

<sup>1</sup>The code provided with BSA was not reproducible, and hence this choice

<sup>2</sup>www.gwern.net/Danbooru2018

Method	BigGAN (128 x 128)		
	Places2.5k	FFHQ2k	CUB6k
	FID ↓	FID ↓	FID ↓
MineGAN	75.50	75.91	69.64
TransferGAN + DISP-Vgg16	162.91 <b>57.35</b>	126.23 <b>44.43</b>	138.87 <b>23.37</b>
FreezeD + DISP-Vgg16	191.04 <b>50.58</b>	161.87 <b>43.90</b>	142.47 <b>26.90</b>
DiffAugment + DISP-Vgg16 + DISP-SimCLR	56.48 30.76 <b>26.65</b>	31.60 23.19 <b>21.06</b>	36.09 15.81 <b>12.36</b>

Table 2: FID of different techniques on limited data image generation. FID (lower is better) is computed using 10k, 7k, 6k generated and real samples (disjoint from training set) for Places2.5k, FFHQ2k, CUB datasets respectively. All above approaches are initialized with BigGAN model pre-trained on ImageNet.

FFHQ [23] datasets. The above choice of datasets follows from the prior work BSA. For methods with pre-training, we finetune SNGAN pre-trained on ImageNet as done in [37] (there is no class label intersection of the above datasets with ImageNet classes). We also show additional results at  $256 \times 256$  resolution on additional datasets (Pandas, Grumpy Cat, Obama) with StyleGAN2 [24] in supplementary.

**Results** Using DISP shows consistent improvement in FID and Recall over all baseline methods as shown in Table 5. Fig 3 shows samples generated via interpolation between conditional embedding of models trained via DISP-Vgg on DiffAugment and vanilla DiffAugment. These results qualitatively show the improvement obtained using our DISP transfer learning approach. Comparatively, the baseline, vanilla DiffAugment, fails to generate realistic interpolation and for the most part, presents memorized training set images. DISP also performs better when training is done from scratch as compared to FreezeD and TransferGAN but is worse than DiffAugment + DISP. We present additional ablation studies in Supplementary.

**Performance on varying number of training images** We vary the number of training examples in Anime dataset from 25-500 for baseline few-shot algorithms and their respective augmentations with DISP-Vgg16. The FID metric comparison in Fig 3c shows the benefit of our approach when used with existing training algorithms. The FID metric for all approaches improves (decreases) with the increase the number of training images with DISP out-performing corresponding baselines. Sample images generated by our approach are shown in Supplementary.

**Memorization Test** To evaluate whether trained GANs are actually generating novel images instead of only memorizing the training set, we calculate FID between images randomly sampled from training set with repetition and the

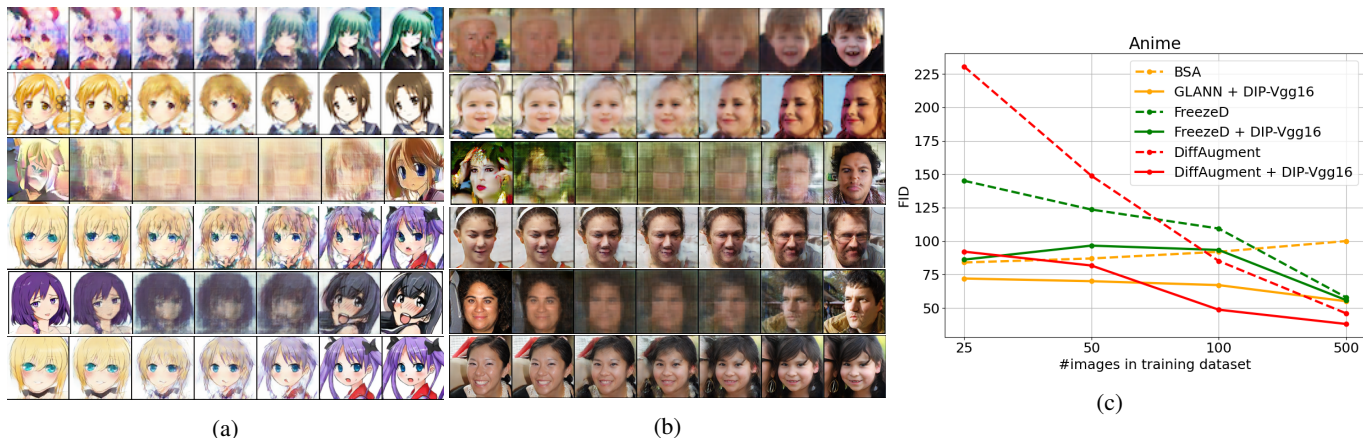


Figure 3: (a) and (b): Sample interpolations between two generated images for models trained in few-shot setting : Scratch (Row 1), Scratch + DISP-Vgg16 (Row 2), FreezeD (Row 3), FreezeD + DISP-Vgg16 (Row 4), DiffAugment (Row 5), DiffAugment + DISP-Vgg16 (Row 6). (c): FID (lower is better) performance graph of few-shot image generation by varying the training samples from 25 to 500 images of Anime dataset for different approaches on SNGAN model.

separate test set for Anime and FFHQ dataset. For Anime dataset, we get an FID of 81.23 and for FFHQ, 100.07. On comparing these numbers to Table 5 we observe that only on using DISP with existing algorithms, we achieve a better FID score suggesting that our approach is able to generate novel/diverse samples instead of just memorizing or overfitting to training data.

**Analyzing the feature space of Vgg-16/SimCLR pre-trained network for Anime dataset** To examine the usefulness of Vgg features on Anime dataset, we evaluate it on the anime character classification task. We took a subset of 70k images from the Anime Face dataset that had labels assigned among the 50 character tags. Each character tag has around 1000-1500 images. We train a single linear classifier on Vgg-16 features of 50k samples and evaluate it on the rest 20k samples. We observe an accuracy of 75% and 67% on training and test sets respectively. When a single linear classifier is trained upon SimCLR features, the respective accuracies were 81% and 63.5%. This highlights that even for fine-grained and out of domain distributions like Anime, pre-trained Vgg-16 features are semantically rich enough to achieve a decent classification score.

## 5.2. Limited Data Image Generation

In many practical scenarios, we have access to moderate number of images (1k-5k) instead of just a few examples, however the limited data may still not be enough to achieve stable GAN training. We show the benefit of our approach in this setting and compare our results with: MineGAN[56], TransferGAN, FreezeD, and DiffAugment. We perform experiments on three  $128 \times 128$  resolution datasets: FFHQ2k, Places2.5k and CUB6k following [56]. FFHQ2k contains 2K training samples from FFHQ [23] dataset. Places2.5k is a subset of Places365 dataset [68] with 500 examples each

sampled from 5 classes (alley, arch, art gallery, auditorium, ball-room). CUB6k is the complete training split of CUB-200 dataset [55]. We use the widely used BigGAN [4] architecture, pre-trained on ImageNet for finetuning. Table 2 shows our results; using DISP consistently improves FID on existing baselines by a significant margin. More implementation details are given in supplementary and sample generated images via our approach are shown in Figure 4.

**Experiments on CIFAR-10 and CIFAR-100** We also experiment with unconditional BigGAN and StyleGAN2 model on CIFAR-10 and CIFAR-100 while varying the amount of data as done in [66]. We compare DISP with DiffAugment on all settings and the results are shown in Table 3. In the limited data setting (5% and 10%) augmenting DiffAugment with DISP gives the best results in terms of FID for both BigGAN and StyleGAN2 architectures. When trained on complete training dataset DISP slightly outperforms DiffAugment on BigGAN architecture. For implementation details, please refer to supplementary.

## 5.3. Large-Scale Image Generation

In order to show the usefulness of our method on large-scale image generation, we carry out experiments on CIFAR-10, CIFAR-100 [26] and ImageNet- $32 \times 32$  datasets with 50k, 50k and  $\sim 1.2M$  training images respectively at  $32 \times 32$  resolution. For a higher  $128 \times 128$  resolution, we perform experiments on FFHQ and LSUN-bedroom [61] datasets with 63k and 3M training samples. We use a ResNet-based architecture for both discriminator and generator similar to BigGAN [4] for all our experiments. We also compare DISP with SSGAN [6] and Self-Conditional GAN [29]. Implementation and training hyperparameter details are provided in Supplementary.

**Results** Table 4 reports the FID, precision and recall score




Figure 4: Samples of generated image in limited data training setting : FreezeD (Row 1), FreezeD + DISP-Vgg16 (Row 2), DiffAugment (Row 3) and DiffAugment + DISP-Vgg16 (Row 4).

Method	CIFAR-10			CIFAR-100		
	100% data	20% data	10% data	100% data	20% data	10% data
BigGAN	17.22	31.25	42.59	20.37	33.25	42.43
+ DISP	9.70	16.24	27.86	12.89	21.70	31.48
+ DiffAugment	10.39	15.12	18.56	13.33	19.78	23.80
+ DiffAugment & DISP	<b>9.52</b>	<b>14.24</b>	<b>18.50</b>	<b>12.70</b>	<b>16.91</b>	<b>20.47</b>
StyleGAN2*	11.07	23.08	36.02	16.54	32.30	45.87
+ DiffAugment*	9.89	12.15	14.5	15.22	16.65	20.75
+ DiffAugment & DISP	<b>9.50</b>	<b>10.92</b>	<b>12.03</b>	<b>14.45</b>	<b>15.52</b>	<b>17.33</b>

Table 3: Comparison of FID on Unconditional CIFAR-10 and CIFAR-100 image generation while varying the amount of training data. Here, all mentioned approaches are trained with random-horizontal flip augmentation of real images. BigGAN-DiffAugment includes consistency regularization [63] following the implementation provided by authors [66]. We report the checkpoint with the best FID value for each model. \* denotes directly reported from paper.

on the generated samples and the test set for baselines and our approach (DISP). For fitting GMM, the number of components are fixed to  $1K$  for all datasets. DISP achieves better FID, precision and recall scores compared to leading baselines. Sample qualitative results and generation with latent interpolation are shown in the supplementary. We also evaluate the quality of inverted images for  $128 \times 128$  resolution on FFHQ and LSUN datasets using Inference via Optimization Measure (IvOM) [33] to emphasize the high instance-level data coverage in the prior space of GANs trained through our approach (details on IvOM calculation are provided in supplementary). Table 7 shows the IvOM and FID metric between inverted and real query images. Figure 5 shows sample inverted images. We observe both from qualitative and quantitative perspective, models trained via DISP inverts a given query image better than the corresponding baselines. We also perform an ablation experiment to analyze the effect of different priors in DISP for CIFAR-100 dataset. As shown in Table 8, the FID metric remains relatively similar for different priors when compared to the baseline.

**Memorization test** For analyzing memorization in GANs, we evaluate it on the recently proposed test to detect data copying [32]. The test calculates whether generated

samples are closer to the training set as compared to a separate test set in the inception feature space using three sample Mann-Whitney U test [30]. The test statistic  $C_T \ll 0$  represents overfitting and data-copying, whereas  $C_T \gg 0$  represents underfitting. We average the test statistic over 5 trials and report the results in Table 6. We can see that using data instance priors during GAN training does not lead to data-copying according to the test statistic except in case of FFHQ dataset where both DISP and baseline  $C_T$  values are also negative.

**Performance gain due to knowledge distillation vs Memorization of real image features** We conduct an additional experiment where we use the features of a Resnet50 network trained on 75% label-corrupted CIFAR-100 as Data Instance Priors to train CIFAR-100 BigGAN architecture. This results in a significantly higher FID (22.82) in comparison to using prior feature from Resnet50-SimCLR trained on clean CIFAR100 dataset, (FID 14.62, Table 8). This highlights that performance depends on the quality of pre-trained network features and not only because features of real images are leveraged as prior during generation. If this was not the case then using features of Resnet50 trained on 75% label-corrupted CIFAR-100 would have resulted in similar performance.

Method	CIFAR-10			CIFAR-100			FFHQ			LSUN-Bedroom			ImageNet32x32		
	FID ↓	P ↑	R ↑	FID ↓	P ↑	R ↑	FID ↓	P ↑	R ↑	FID ↓	P ↑	R ↑	FID ↓	P ↑	R ↑
Baseline	19.73	0.64	<b>0.70</b>	24.66	0.61	<b>0.67</b>	21.67	0.77	0.47	9.89	0.58	0.42	16.19	0.60	<b>0.67</b>
SSGAN	15.65	0.67	0.68	21.02	0.61	0.65	-	-	-	7.68	0.59	0.50	17.18	0.61	0.65
Self-Cond GAN	16.72	0.71	0.64	21.8	0.64	0.60	-	-	-	-	-	-	15.56	<b>0.66</b>	0.63
DISP-Vgg16	<b>11.24</b>	<b>0.74</b>	0.64	<b>15.71</b>	<b>0.70</b>	0.62	<b>15.83</b>	0.76	<b>0.55</b>	4.99	<b>0.66</b>	<b>0.54</b>	<b>12.11</b>	0.64	0.62
DISP-SimCLR	14.42	0.68	0.65	20.08	0.67	0.62	16.62	<b>0.77</b>	0.53	<b>4.92</b>	0.62	0.53	14.99	0.60	0.63

Table 4: Comparison of DISP with Baseline, SSGAN[6] and Self-Cond GAN[29] in large-scale image generation setting on FID, Precision and Recall metrics.

Method	50k	100k	200k	500k
GMM	4.99	4.92	4.81	4.43
Time (s)	383.96	1063.99	1993.93	4397.56

Table 5: Relationship among the number of random samples used in the GMM, FID value obtained and the time taken for learning the GMM.

Methods	CIFAR-10	CIFAR-100	FFHQ	LSUN	ImageNet 32x32
	$C_T$	$C_T$	$C_T$	$C_T$	$C_T$
Baseline	3.02	4.26	-0.15	2.59	10.5
DISP-Vgg16	1.58	3.05	-0.81	1.06	8.53
DISP-SimCLR	2.86	3.48	-1.49	0.13	9.91

Table 6: Test for evaluating data-copy and memorization in GANs [32] for different approaches and datasets. Test statistic  $C_T \ll 0$  denotes overfitting and data-copying, and  $C_T \gg 0$  represents under-fitting.

Method	FFHQ		LSUN-Bedroom	
	IvOM ↓	FID ↓	IvOM ↓	FID ↓
Baseline	0.0386	85.06	0.0517	<b>115.02</b>
+ DISP-Vgg16	0.0142	73.85	0.0191	129.4
+ DISP-SimCLR	<b>0.0125</b>	<b>71.44</b>	<b>0.0161</b>	116.11

Table 7: IvOM and FID measure on 500 random test images of FFHQ and LSUN-Bedroom datasets.

**Relationship between the number of random samples used for fitting GMM and its corresponding FID** Fitting a GMM model is inhibitive for large-scale datasets e.g. ImageNet and LSUN-Bedroom where the training data is in millions. We observed this during our experiments and therefore use a subset of randomly sampled 200K instances for these datasets for fitting GMM in Table 4. Table 5 further shows the relationship between the number of random samples used for fitting GMM and the corresponding FID (average of 3 runs with a standard deviation of less than 1%) on the LSUN-Bedroom dataset for DISP-Vgg16 trained model on LSUN-Bedroom. As can be seen, even a small subset of training data can still be used to achieve better performance than baselines in relatively less time. This experiment was performed on a system with 32 CPU cores, 64 GB RAM, and processor Intel(R) Xeon(R) CPU @ 2.20GHz.




Figure 5: Images generated through IvOM for randomly sampled test set images on FFHQ and LSUN-Bedroom. (Top to Bottom:) Original images, Baseline, Baseline + DISP-Vgg16, Baseline + DISP-SimCLR.

Method	CIFAR-100
Baseline	24.66
+ DISP-SimCLR (ImageNet)	16.26
+ DISP-SimCLR (CIFAR-100)	14.62
+ DISP-ResNet50 (Places-365)	14.68
+ DISP-Resnet50 (ImageNet)	14.62

Table 8: Comparison of FID when using prior from different pre-trained models on CIFAR-100.

## 6. Conclusion

In this work, we present a novel instance level prior based transfer learning approach to improve the quality and diversity of images generated using GANs when a few training data samples are available. By leveraging features as priors from rich source domain in limited unsupervised image synthesis, we show the utility of our simple yet effective approach on various standard vision datasets and GAN architectures. We demonstrate the efficacy of our approach in image generation with limited data, where it achieves the new state-of-the performance, as well as on large-scale settings. As future work, it would be interesting to explore the application of prior information in image editing tasks.



## References

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223, 2017.
- [2] Sanjeev Arora, Rong Ge, Yingyu Liang, Tengyu Ma, and Yi Zhang. Generalization and equilibrium in generative adversarial nets (gans). In *ICML*, 2017.
- [3] Piotr Bojanowski, Armand Joulin, David Lopez-Pas, and Arthur Szlam. Optimizing the latent space of generative networks. In *International Conference on Machine Learning*, pages 600–609, 2018.
- [4] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2018.
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [6] Ting Chen, Xiaohua Zhai, Marvin Ritter, Mario Lucic, and Neil Houlsby. Self-supervised gans via auxiliary rotation loss. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12154–12163, 2019.
- [7] Harm De Vries, Florian Strub, Jérémie Mary, Hugo Larochelle, Olivier Pietquin, and Aaron C Courville. Modulating early visual processing by language. In *Advances in Neural Information Processing Systems*, pages 6594–6604, 2017.
- [8] Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In *Advances in neural information processing systems*, pages 1486–1494, 2015.
- [9] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *International conference on machine learning*, pages 647–655, 2014.
- [10] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. A learned representation for artistic style. *arXiv preprint arXiv:1610.07629*, 2016.
- [11] Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT Press, 2016.
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [13] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777, 2017.
- [14] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017.
- [15] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015.
- [16] Yedid Hoshen, Ke Li, and Jitendra Malik. Non-adversarial image synthesis with generative latent nearest neighbors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5811–5819, 2019.
- [17] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017.
- [18] Ali Jahanian, Lucy Chai, and Phillip Isola. On the “steerability” of generative adversarial networks. In *ICLR*, 2020.
- [19] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [20] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. In *ICLR*, 2018.
- [21] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *arXiv preprint arXiv:2006.06676*, 2020.
- [22] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 12104–12114. Curran Associates, Inc., 2020.
- [23] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4401–4410, 2019.
- [24] 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.
- [25] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014.
- [26] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10. *URL* <http://www.cs.toronto.edu/kriz/cifar.html>, 2010.
- [27] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. In *Advances in Neural Information Processing Systems*, pages 3927–3936, 2019.
- [28] Ke Li and Jitendra Malik. Implicit maximum likelihood estimation. *arXiv preprint arXiv:1809.09087*, 2018.
- [29] Steven Liu, Tongzhou Wang, David Bau, Jun-Yan Zhu, and Antonio Torralba. Diverse image generation via self-conditioned gans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [30] Henry B Mann and Donald R Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, pages 50–60, 1947.

- [31] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2794–2802, 2017.
- [32] Casey Meehan, Kamalika Chaudhuri, and Sanjoy Dasgupta. A non-parametric test to detect data-copying in generative models. *arXiv preprint arXiv:2004.05675*, 2020.
- [33] Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. Unrolled generative adversarial networks. *ICLR*, 2017.
- [34] Takeru Miyato, Toshiaki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *ICLR*, 2018.
- [35] Takeru Miyato and Masanori Koyama. cgans with projection discriminator. In *International Conference on Learning Representations*, 2018.
- [36] 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.
- [37] Atsuhiko Noguchi and Tatsuya Harada. Image generation from small datasets via batch statistics adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2750–2758, 2019.
- [38] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *International conference on machine learning*, pages 2642–2651, 2017.
- [39] Maxime Oquab, Leon Bottou, Ivan Laptev, and Josef Sivic. Learning and transferring mid-level image representations using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014.
- [40] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.
- [41] Antoine Plumerault, Hervé Le Borgne, and Céline Hudelot. Controlling generative models with continuous factors of variations. In *ICLR*, 2020.
- [42] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *International Conference on Learning Representations*, 2016.
- [43] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *ICML*, 2014.
- [44] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In *ICLR*, 2015.
- [45] Alexander Sage, Eirikur Agustsson, Radu Timofte, and Luc Van Gool. Logo synthesis and manipulation with clustered generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5879–5888, 2018.
- [46] Mehdi S. M. Sajjadi, Olivier Bachem, Mario Lučić, Olivier Bousquet, and Sylvain Gelly. Assessing Generative Models via Precision and Recall. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [47] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016.
- [48] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *CVPR*, 2020.
- [49] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [50] Dustin Tran, Rajesh Ranganath, and David M Blei. Deep and hierarchical implicit models. *arXiv preprint arXiv:1702.08896*, 7:3, 2017.
- [51] Michael Tschannen, Eirikur Agustsson, and Mario Lucic. Deep generative models for distribution-preserving lossy compression. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [52] Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4068–4076, 2015.
- [53] Aaron Van den Oord, Nal Kalchbrenner, Lasse Espeholt, Oriol Vinyals, Alex Graves, et al. Conditional image generation with pixelcnn decoders. In *Advances in neural information processing systems*, pages 4790–4798, 2016.
- [54] Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International Conference on Machine Learning*, pages 1747–1756, 2016.
- [55] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- [56] Yaxing Wang, Abel Gonzalez-Garcia, David Berga, Luis Herranz, Fahad Shahbaz Khan, and Joost van de Weijer. Minegan: Effective knowledge transfer from gans to target domains with few images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [57] Yaxing Wang, Chenshen Wu, Luis Herranz, Joost van de Weijer, Abel Gonzalez-Garcia, and Bogdan Raducanu. Transferring gans: generating images from limited data. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 218–234, 2018.
- [58] Lei Xu and Michael I. Jordan. On convergence properties of the em algorithm for gaussian mixtures. *Neural Computation*, 8(1):129–151, Jan. 1996.
- [59] Xin Yi, Ekta Walia, and Paul Babyn. Generative adversarial network in medical imaging: A review. *Medical image analysis*, 58:101552, 2019.
- [60] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In *Advances in Neural Information Processing Systems 27*, pages 3320–3328, 2014.
- [61] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, 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.

- [62] Ning Yu, Ke Li, Peng Zhou, Jitendra Malik, Larry Davis, and Mario Fritz. Inclusive gan: Improving data and minority coverage in generative models. In *European Conference on Computer Vision*, pages 377–393. Springer, 2020.
- [63] Han Zhang, Zizhao Zhang, Augustus Odena, and Honglak Lee. Consistency regularization for generative adversarial networks. In *International Conference on Learning Representations*, 2019.
- [64] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [65] Miaoyun Zhao, Yulai Cong, and Lawrence Carin. On leveraging pretrained gans for generation with limited data. In *ICML*, 2020.
- [66] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable augmentation for data-efficient gan training. *arXiv preprint arXiv:2006.10738*, 2020.
- [67] Zhengli Zhao, Sameer Singh, Honglak Lee, Zizhao Zhang, Augustus Odena, and Han Zhang. Improved consistency regularization for gans. *arXiv preprint arXiv:2002.04724*, 2020.
- [68] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Object detectors emerge in deep scene cnns. *arXiv preprint arXiv:1412.6856*, 2014.
- [69] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *Proceedings of European Conference on Computer Vision (ECCV)*, 2020.
- [70] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networkss. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017.