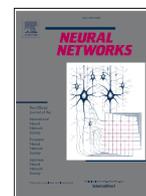




ELSEVIER

Contents lists available at ScienceDirect

## Neural Networks

journal homepage: [www.elsevier.com/locate/neunet](http://www.elsevier.com/locate/neunet)

Full Length Article

# HP-GAN: Harnessing pretrained networks for GAN improvement with FakeTwins and discriminator consistency

 Geonhui Son <sup>a</sup>, Jeong Ryong Lee <sup>a</sup>, Dosik Hwang <sup>a,b,c,d,\*</sup>
<sup>a</sup> School of Electrical and Electronic Engineering, Yonsei University, Seoul, Republic of Korea

<sup>b</sup> Department of Radiology and Research Institute of Radiological Science and Center for Clinical Imaging Data Science, Yonsei University College of Medicine, Seoul, Republic of Korea

<sup>c</sup> Artificial Intelligence and Robotics Institute, Korea Institute of Science and Technology, Seoul, Republic of Korea

<sup>d</sup> Department of Oral and Maxillofacial Radiology, Yonsei University College of Dentistry, Seoul, Republic of Korea

## ARTICLE INFO

## Keywords:

 Image generation  
 Generative adversarial network  
 Pretrained network  
 Self-supervised learning

## ABSTRACT

Generative Adversarial Networks (GANs) have made significant progress in enhancing the quality of image synthesis. Recent methods frequently leverage pretrained networks to calculate perceptual losses or utilize pretrained feature spaces. In this paper, we extend the capabilities of pretrained networks by incorporating innovative self-supervised learning techniques and enforcing consistency between discriminators during GAN training. Our proposed method, named HP-GAN, effectively exploits neural network priors through two primary strategies: FakeTwins and discriminator consistency. FakeTwins leverages pretrained networks as encoders to compute a self-supervised loss and applies this through the generated images to train the generator, thereby enabling the generation of more diverse and high quality images. Additionally, we introduce a consistency mechanism between discriminators that evaluate feature maps extracted from Convolutional Neural Network (CNN) and Vision Transformer (ViT) feature networks. Discriminator consistency promotes coherent learning among discriminators and enhances training robustness by aligning their assessments of image quality. Our extensive evaluation across seventeen datasets-including scenarios with large, small, and limited data, and covering a variety of image domains-demonstrates that HP-GAN consistently outperforms current state-of-the-art methods in terms of Fréchet Inception Distance (FID), achieving significant improvements in image diversity and quality. Code is available at: <https://github.com/higun2/HP-GAN>.

## 1. Introduction

Generative Adversarial Networks (GANs) (Goodfellow et al., 2020) are a prominent method in generative modeling, capable of synthesizing high quality, photorealistic images. This capability is achieved through a min-max game between two components: a generator, which creates synthetic data, and a discriminator, which distinguishes between synthetic and real data. Proper training of these two components is crucial for achieving high quality synthesis, as the discriminator acts as an adaptive loss function for the generator.

Nevertheless, GANs still encounter significant challenges such as non-convergence, training instability, and mode collapse (Arjovsky & Bottou, 2017; Mescheder et al., 2018; Zhang & Khoreva, 2019). These limitations are particularly pronounced when computational resources and datasets are limited, as is often the case in medical imaging of

rare diseases, specific sets of celebrity portraits, or the artwork of a particular artist. Previous studies have attempted to mitigate these issues through various approaches, including modifications in network architectures (Karras et al., 2019; Liu et al., 2021; Radford et al., 2016; Sauer et al., 2021; Zhang et al., 2019), refinement of objective functions (Arjovsky et al., 2017; Bellemare et al., 2017; Deshpande et al., 2018; Kunkel & Trabs, 2025; Li et al., 2017; Nowozin et al., 2016), regularization of weights and gradients (Arjovsky et al., 2017; Fedus et al., 2018; Mescheder et al., 2018; Miyato et al., 2018; Roth et al., 2017; Salimans et al., 2016), and use of side information (Kim et al., 2023, 2021; Wang et al., 2018a; Zhang et al., 2020a, 2017).

Data augmentation, conventionally employed to mitigate overfitting in deep neural networks (Cubuk et al., 2018, 2020; Zhang et al., 2018a), has also been adapted to enhance GAN training. Recent studies indicate that augmentation applied to both real and synthesized samples

\* Corresponding author.

E-mail addresses: [higun2@yonsei.ac.kr](mailto:higun2@yonsei.ac.kr) (G. Son), [emaker@yonsei.ac.kr](mailto:emaker@yonsei.ac.kr) (J.R. Lee), [dosik.hwang@yonsei.ac.kr](mailto:dosik.hwang@yonsei.ac.kr) (D. Hwang).

<https://doi.org/10.1016/j.neunet.2026.108666>

Received 18 June 2025; Received in revised form 5 December 2025; Accepted 27 January 2026

Available online 29 January 2026

0893-6080/© 2026 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

improves synthesis performance. Karras et al. (2020a), Tran et al. (2021), Zhang and Khoreva (2019), Zhang et al. (2020b), Zhao et al. (2020b,c). However, the learning objective of the discriminator remains to distinguish between the real and fake images, which degrade performance when training data is limited (Yang et al., 2021).

While transfer learning offers a potential solution (Ham et al., 2020; Mo et al., 2020; Wang et al., 2020, 2018b; Zhao et al., 2020a), pretrained networks are not always compatible with the given training dataset, and improper fine-tuning can result in performance degradation (Zhao et al., 2020b). Pretrained networks have extensive applications in image-to-image translation tasks, serving as perceptual loss functions (Dosovitskiy & Brox, 2016; Johnson et al., 2016) or perceptual discriminators (Richter et al., 2023; Sungatullina et al., 2018). Projected GAN (Sauer et al., 2021) utilizes pretrained network to project generated and real samples into pretrained feature spaces, leading to significant improvements in performance, training stability, time, and data efficiency. StyleGAN-XL (Sauer et al., 2022) utilizes both Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs) networks for complementary effects. Furthermore, Vision-aided GAN (Kumari et al., 2022) employs multiple networks by gradually increasing the number of pretrained networks during training.

In this paper, we propose HP-GAN, a novel approach that leverages pretrained networks to improve the performance of the generator by incorporating self-supervised learning (SSL) techniques. First, we introduce FakeTwins, which utilizes pretrained feature networks as SSL encoders and trains the generator with SSL based on Barlow Twins (Zbontar et al., 2021) through the generated images. Recently, contrastive learning based SSL methods, such as SimCLR (Chen et al., 2020a) and MoCov2 (Chen et al., 2020b), have been integrated in GANs (Jeong & Shin, 2021; Li et al., 2022; Yang et al., 2021). These methods incorporate contrastive learning as an auxiliary task during GAN training, thereby enhancing the discriminative capabilities of the discriminator. In contrast, we leverage information maximization based methods (Bardes et al., 2022; Ermolov et al., 2021; Zbontar et al., 2021) to enhance the capabilities of the generator in different way. Barlow Twins is an SSL method that minimizes redundancy between the output features of two identical networks processing differently distorted (or augmented) versions of the same sample. We assume that Barlow Twins with a pretrained network yields lower loss values for batches with diverse images, motivating the development of FakeTwins to enhance the diversity and quality of generated images.

Furthermore, we introduce the discriminator consistency loss, a regularization technique designed to ensure cohesive learning among multiple discriminators. We employ both Convolutional Neural Network (CNN) and Vision Transformer (ViT) as feature networks (Sauer et al., 2022). However, due to the inherent architectural differences between CNNs and ViTs, the discriminator outputs from these networks do not align identically. The discriminator consistency loss mitigates these discrepancies by encouraging a high level of consensus among the discriminators regarding the quality of images. This approach ensures that the generator receives uniform and coherent feedback, leading to more stable and robust training. In other words, the alignment of discriminator outputs leverages the structural differences between the pretrained CNN and ViT to provide a comprehensive assessment of the generated image.

To assess the quality of images produced by generative models, we employ several metrics for comprehensive comparison. Fréchet Inception Distance (FID) (Heusel et al., 2017) measures the disparity in distribution density within the feature space, providing a quantitative evaluation of image quality. Additionally, Kernel Inception Distance (KID) (Bińkowski et al., 2018) is an unbiased alternative to FID, often yielding more stable evaluations for smaller sample sizes. Precision and recall (Kynkäänniemi et al., 2019) offer additional insight by quantifying the proportion of generated images that resemble the training set and the model's ability to replicate the training set, respectively.

The main contributions of our work are summarized as follows:

1. We propose the discriminator consistency loss, a novel regularization technique that mitigates disparities in discriminator outputs. Discriminator consistency improves training stability and overall performance by providing coherent feedback to the generator, ensuring robust and stable training.
2. We introduce FakeTwins, an innovative approach that incorporates self-supervised learning based method to train the generator. By utilizing pretrained feature networks as SSL encoders, FakeTwins enhances the diversity and fidelity of the generated images.
3. Through extensive experimental evaluations on seventeen diverse image-domain datasets, HP-GAN consistently outperforms existing state-of-the-art methods in generative modeling. Our method demonstrates significant improvements in image synthesis quality, achieving state-of-the-art performance across various benchmarks.

## 2. Related works

### 2.1. Generative adversarial networks

Generative Adversarial Networks (GANs) (Goodfellow et al., 2020) model the data distribution of a given training dataset. This is accomplished by setting up a min-max game involving two neural networks known as the generator and the discriminator. The generator  $G$  maps latent code  $z$  sampled from the latent space  $\mathcal{Z}$ , and produces realistic-looking samples  $G(z)$ . The discriminator  $D$  is fed either real samples  $x$  or synthetic samples  $G(z)$  from the generator and tries to accurately distinguish them as real or fake. The min-max objective function of GAN is as follows:

$$\min_G \max_D (\mathbb{E}_x [\log D(x)] + \mathbb{E}_z [\log(1 - D(G(z)))]). \quad (1)$$

### 2.2. Data augmentation in GANs

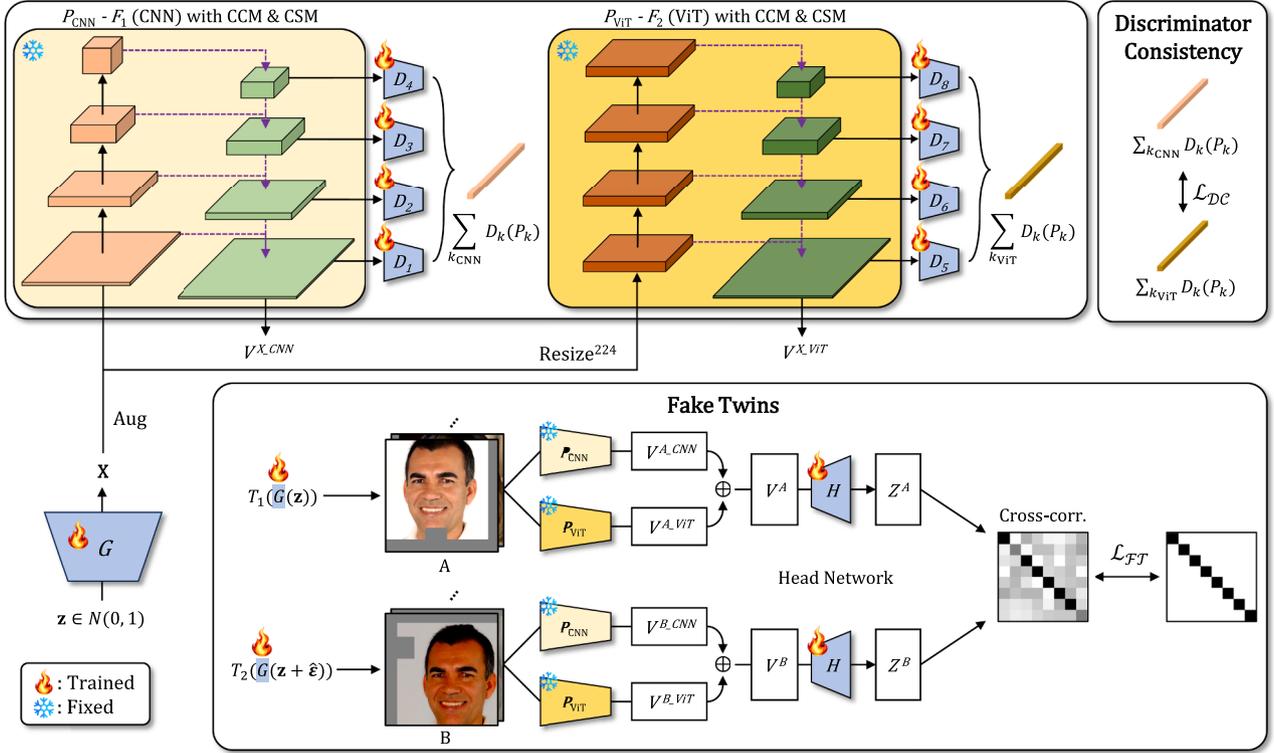
Data augmentation maximizes the utilization of datasets to improve performance on tasks such as classification (Cubuk et al., 2018, 2020; Zhang et al., 2018a) by artificially increasing the size and diversity of training datasets. In GANs, data augmentation also has been employed to prevent overfitting, stabilize training, and improve the quality of synthesized images, especially when a limited number of training samples are available (Karras et al., 2020a; Tran et al., 2021; Zhang & Khoreva, 2019; Zhang et al., 2020b; Zhao et al., 2020b,c).

For instance, CR-GAN (Zhang et al., 2020b) augments the real images and introduces a consistency loss for training the discriminator. The other methods, ADA (Karras et al., 2020a) and DA (Zhao et al., 2020b), propose a concept of applying differential augmentation on both real and synthetic images to address the insufficiency of the training data and thus enhance the quality and diversity of the generated data. ADA (Karras et al., 2020a) utilizes data augmentation through an adaptive technique. The strength of the augmentation is adjusted according to the current level of training stability, thus maintaining a balance between training stability and the diversity of the generated images. DA (Zhao et al., 2020c) introduces a differentiable augmentation strategy. This allows the gradients from the discriminator to flow back to the generator through the augmentations.

In a different direction, DANI (Zhang et al., 2024) introduces feature-space augmentation by injecting learnable noise into both the generator and discriminator. Instead of transforming input images, DANI simulates the regularization effect of augmentation through adaptive noise, improving stability and performance in low-data regimes.

### 2.3. Self-supervised learning

Self-supervised learning (SSL) is an emerging paradigm that aims to learn meaningful representations from unlabeled data. SSL has demonstrated significant success in various domains, including image classification, object detection, and natural language processing



**Fig. 1.** Algorithmic overview. The images are augmented and resized, then passed through pretrained CNN and ViT networks, along with their respective feature mixing blocks (CCM + CSM, dashed purple arrows). The processed multi-scale feature maps are then subjected to independent discriminators. During training, we impose consistency between discriminator outputs from different feature networks ( $\mathcal{L}_{DC}$ ). Lastly, we propose the FakeTwins loss ( $\mathcal{L}_{FT}$ ), which is applied through fake images, with the aim of maximizing the diversity of images. FakeTwins is an SSL based approach which utilizes pretrained networks as encoders.  $\oplus$  indicates concatenation.

(Bachman et al., 2019; Bardes et al., 2022; Caron et al., 2020; Chen et al., 2020a,b; Chen & He, 2021; Gidaris et al., 2021; Grill et al., 2020; He et al., 2020; Misra & van der, 2020; Tian et al., 2020; Zbontar et al., 2021) SSL methods typically learn representations that remain invariant under various transformations or different distortions.

In contrastive learning (Chen et al., 2020a,b; van den et al., 2018; He et al., 2020; Misra & van der, 2020), the model learns an association between a positive pair (two augmented versions of the same image) and distinguishes it from a set of negative samples (augmented versions of different images). Contrastive methods have shown effectiveness, they also have limitations, such as the difficulty of selecting negative samples and the requirement of large batch sizes.

Information maximization methods (Bardes et al., 2022; Ermolov et al., 2021; Lee et al., 2025; Zbontar et al., 2021) aim to maximize the information content of the embedding. These methods decorrelate each pair of embedding variables that contain redundant information, leading to the redundancy reduction that enhances the information content of the embedding vectors. Moreover, information maximization methods do not require negative samples, asymmetric networks, and large batch sizes.

Barlow Twins (Zbontar et al., 2021) forces the normalized cross-correlation matrix of the two embeddings towards the identity. More specifically, the procedure starts by generating two distorted views for images of a batch  $X$  sampled from a dataset. Two batches of views  $A$  and  $B$  are created by different data augmentations ( $T_1$  and  $T_2$ ), and then encoded into representations  $V^A$  and  $V^B$  by the encoder network  $f$ . The representations are fed to the linear head network  $\phi$ , which produces the embeddings  $Z^A$  and  $Z^B$ . The loss function is constructed to encourage decorrelation between the feature representations of the two augmented views, and mathematically represented

as:

$$A = T_1(X), \quad V^A = f(A), \quad Z^A = \phi(V^A), \quad (2)$$

$$B = T_2(X), \quad V^B = f(B), \quad Z^B = \phi(V^B), \quad (3)$$

$$\mathcal{L}_{BT}(Z^A, Z^B) \triangleq \sum_i (1 - C_{ii})^2 + \lambda_1 \sum_i \sum_{j \neq i} C_{ij}^2, \quad (4)$$

where  $\lambda_1$  is a constant that balances the importance of the first and second terms of the loss, and is set to 0.005.  $C$  is the cross-correlation matrix between the outputs of the two identical networks, computed across the batch dimension:

$$C_{ij} \triangleq \frac{\sum_b z_{b,i}^A z_{b,j}^B}{\sqrt{\sum_b (z_{b,i}^A)^2} \sqrt{\sum_b (z_{b,j}^B)^2}}, \quad (5)$$

where  $b$  indexes batch samples and  $i, j$  index the vector dimension of the network's output.  $C$  is a square matrix with the dimensions of the network's output, and its values range between -1 (perfect anti-correlation) and 1 (perfect correlation). As a result, the embedding vectors of distorted versions of a sample become similar, while simultaneously reducing the redundancy within vector components.

#### 2.4. Self-supervised learning in GANs

The incorporation of SSL in GANs is a recent area of research. Tran et al. (2019) propose a multi-class minimax game where a rotation-based classifier distinguishes between transformed real and fake images, encouraging the generator to align with the self-supervised task. Hou et al. (2021) introduce label augmentation by integrating transformation labels into the adversarial objective, unifying self-supervised and GAN training through a single multi-class classification framework.

**Table 1**

Ablation study on FFHQ and Pokemon dataset. Results for different methods and configurations. Perceptual path length (PPL) (Karras et al., 2019) is computed in  $\mathcal{Z}$  only when  $\mathbf{z} \in \mathbb{R}^{64}$ , without the central crop. The best results are in bold.

Configuration	FFHQ (70k)				Pokemon (833)			
	FID ↓	PPL (full / end) ↓	Pr ↑	Rec ↑	FID ↓	PPL (full / end) ↓	Pr ↑	Rec ↑
<b>A</b> FastGAN (Liu et al., 2021)	12.69	-	<b>0.716</b>	0.184	81.86	-	0.731	0.004
<b>B</b> + Projected GAN (CNN as $F_1$ ) (Sauer et al., 2021)	3.39	-	0.654	0.464	26.36	-	<b>0.809</b>	0.259
<b>C</b> + ViT as $F_2$ & small $\mathbf{z}$ & blur reg. (Sauer et al., 2022)	2.29	150.3 / 150.1	0.672	0.497	24.70	568.5 / 574.3	0.793	0.164
<b>D</b> + Discriminator consistency	1.86	134.3 / 133.4	0.664	0.524	24.02	548.9 / 554.8	0.788	0.210
<b>E</b> + FakeTwins (HP-GAN)	<b>1.69</b>	<b>132.50 / 131.9</b>	0.642	<b>0.545</b>	<b>23.62</b>	<b>450.6 / 451.3</b>	0.768	<b>0.310</b>

Recent approaches involve applying contrastive learning to GANs (Gou et al., 2024; Jeong & Shin, 2021; Li et al., 2022; Wang et al., 2024; Yang et al., 2021) and demonstrate that SSL loss improves the generation performance. ContraD (Jeong & Shin, 2021) reformulates the conventional task to distinguish real and fake samples with two different contrastive losses: the SimCLR loss (Chen et al., 2020a) on the real samples and the supervised contrastive loss (Khosla et al., 2020) on the fake samples. InsGen (Yang et al., 2021) introduces a self-supervised auxiliary task for the discriminator, aiming to enhance its discriminative power by following MoCo-v2 (Chen et al., 2020b). By distinguishing images via contrastive learning, synthesis performance is enhanced. FakeCLR (Li et al., 2022) improves InsGen by introducing new strategies such as applying contrastive loss only to the fake images and implementing noise-related latent augmentation. The implementation of contrastive loss also indirectly changes the generator’s training objectives. As a result, the discriminator and generator are trained better with information from instance discrimination, resulting in improved synthesis quality and diversity.

And for the conditional image generation, IG-GAN (Casanova et al., 2021) trains GANs conditioned on embeddings on SwAV (Caron et al., 2020). ReACGAN (Kang et al., 2021) and ContraGAN (Kang & Park, 2020) introduce contrastive loss into discriminator training, enabling the discriminator to compare visual differences between multiple images and learn the fine-grained representation. NoisyTwins (Rangwani et al., 2023) introduces a self-supervised decorrelation loss inspired by Barlow Twins, applied in the latent space of class embeddings. However, this method inherently relies on the class labels and operates only in conditional settings. In contrast, our method applies the loss directly to image features extracted from high-capacity pretrained networks, enabling class-agnostic learning. This allows richer, perceptually-aligned redundancy reduction that guides the generator more effectively across a broader range of scenarios.

Unlike previous methods (Casanova et al., 2021; Jeong & Shin, 2021; Kang & Park, 2020; Kang et al., 2021; Li et al., 2022; Yang et al., 2021), which typically treat the discriminator’s feature extractor as the trainable SSL encoder and mainly use contrastive, instance-discrimination losses, HP-GAN uses fixed pretrained networks as encoders and applies a negative-free, information-maximization SSL objective directly to generator outputs. Thus, FakeTwins directly regularizes the generator in a pretrained feature space, rather than only improving the discriminator’s discriminative power.

## 2.5. Pretrained models in GANs

Pretrained models, especially those trained with large-scale supervised, unsupervised or self-supervised learning (Caron et al., 2020; Chen et al., 2020a; He et al., 2016; Krizhevsky et al., 2017; Radford et al., 2021), have shown impressive performance in a wide range of tasks through the use of useful representations, transfer learning, and fine-tuning (Donahue et al., 2014; Huh et al., 2016; Kornblith et al., 2019; Oquab et al., 2014; Saenko et al., 2010; Yosinski et al., 2014; Zeiler & Fergus, 2014). Pretrained models have proven effective in reducing the need for extensive training datasets and training time.

Recent works have proposed the idea of leveraging pretrained representations for GANs, thus alleviating the burden of learning from scratch and accelerating convergence speed (Ham et al., 2020; Mo et al., 2020; Wang et al., 2020, 2018b; Zhao et al., 2020a). Another methods introduce perceptual loss function (Dosovitskiy & Brox, 2016; Johnson et al., 2016), which is used for conditional image synthesis tasks (Chen & Koltun, 2017; Gatys et al., 2016; Ledig et al., 2017; Park et al., 2019; Wang et al., 2018a). By comparing the feature representations of the real and generated images in the pretrained model’s feature space, they guide the generator to generate images that are perceptually close to the real images. More recent methods such as DreamSim (Fu et al., 2023) learn human-aligned perceptual similarity metrics from synthetic data and large pretrained feature spaces, providing stronger feature-based distances for analyzing GAN representations. Image-to-image translation is performed by combining perceptual loss and adversarial loss. In addition, the perceptual discriminator (Richter et al., 2023; Sungatullina et al., 2018) that utilizes a pretrained VGG (Simonyan & Zisserman, 2015) is proposed.

On the other hand, pretrained models are used in different ways for unconditional GAN. Projected GAN (Sauer et al., 2021) proposes the concept of projecting generated and real samples into a pretrained feature space. The Projected GAN objective is formulated as follows:

$$\min_G \max_{D_k} \sum_k \left( \mathbb{E}_{\mathbf{x}} [\log D_k(P_k(\mathbf{x}))] + \mathbb{E}_{\mathbf{z}} [\log (1 - D_k(P_k(G(\mathbf{z}))))] \right), \quad (6)$$

where  $P_k$  is a set of feature projectors and  $D_k$  is a set of independent discriminators operating on different feature projections. The feature projectors comprise a pretrained feature network  $F$ , along with cross-channel mixing (CCM) and cross-scale mixing (CSM) layers. The aim of CCM and CSM is to dilute prominent features, preventing the discriminators from focusing on a certain subset of its input feature space. Both layers utilize differentiable random projections that are randomly initialized and not trained during GAN training. CCM mixes features across channels using random  $1 \times 1$  convolutions, and CSM mixes features across scales using residual random  $3 \times 3$  convolutions and bilinear upsampling. The output of CSM is a feature pyramid with four feature maps at different resolutions, and four discriminators operate independently on these feature maps. StyleGAN-XL (Sauer et al., 2022) introduces a progressive growing strategy (Karras et al., 2018) to StyleGAN3 (Karras et al., 2021), demonstrating that incorporating multiple pretrained feature extractors improves generation quality. Vision-aided GAN (Kumari et al., 2022) also validates the effectiveness of ensembling multiple pretrained networks.

However, these methods either use a single feature backbone or treat multiple backbones independently, with each discriminator separately distinguishing between real and fake images. In contrast, our method goes further by leveraging pretrained networks not only for real/fake discrimination but also as encoders for self-supervised learning. This dual role of pretrained networks provides richer supervisory signals and promotes the generation of more diverse and high-quality images. Moreover, our method introduces an explicit consistency regularization term that aligns the outputs of heterogeneous discriminators, encouraging stable learning dynamics and coherent generator updates, which is conceptually related to consistency learning (Guo et al., 2025).

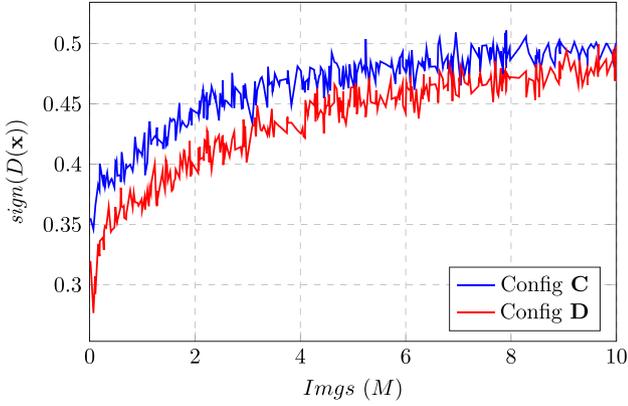


Fig. 2. The signed real logits of the discriminator  $\text{sign}(D(\mathbf{x}))$  according to Config D and E on the FFHQ dataset.  $\text{Imgs}(M)$  denotes the cumulative number of training images (in millions).

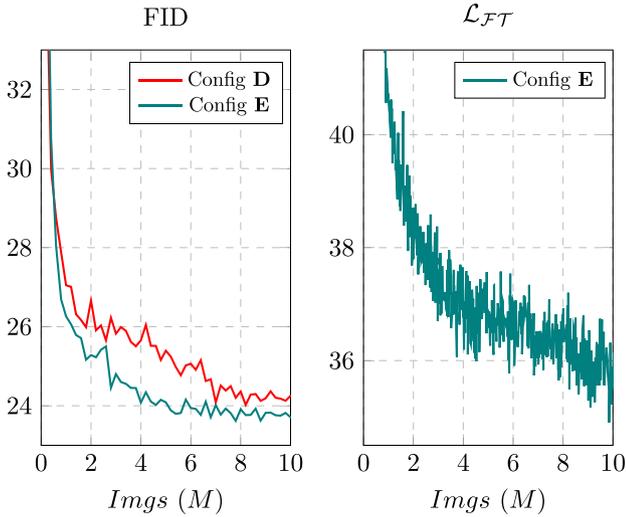


Fig. 3. FID and  $\mathcal{L}_{FT}$  according to Config D and E on the pokemon dataset.  $\text{Imgs}(M)$  denotes the cumulative number of training images (in millions).

### 3. Methodology

In this section, we present our novel approach, which combines the capabilities of pretrained models with self-supervised learning (SSL) to produce diverse and high-quality synthetic images. Our proposed method, termed Harnessing Pretrained Networks GAN (HP-GAN), builds on the FastGAN baseline (Config A), incorporating a series of progressive enhancements to improve performance and robustness. The step-by-step introduction of these modifications is detailed below, with their respective impacts summarized in Table 1.

First, we restructure the traditional single discriminator setup into a multi-discriminator architecture, like Projected GAN (Config B). Next, taking cues from StyleGAN-XL, we introduce a hybrid architecture that integrates CNNs and ViT as feature networks (Config C).

The subsequent enhancement involves the application of discriminator consistency (Config D). By ensuring that the outputs of discriminators, which are based on different feature networks (CNN and ViT), are aligned, we promote a higher level of consensus on image quality, leading to more uniform feedback to the generator and stable training. Finally, we incorporate FakeTwins, resulting in Config E. By leveraging pretrained networks, we enhance the generator’s ability to produce diverse images. Our integrated framework, as depicted in Fig. 1, demonstrates how these incremental innovations collectively contribute to significant improvements in GAN performance.

#### 3.1. Training GANs

Our proposed method, HP-GAN, utilizes feature projectors with multiple feature networks and multi-scale discriminators. We employ a hinge loss (Lim & Ye, 2017) to train our  $D$  and  $G$ . Concretely, loss functions are formalized as:

$$\mathcal{L}_D = \sum_k \left( \mathbb{E}_{\mathbf{x}} [\max(0, 1 - D_k(P_k(\mathbf{x}))) + \mathbb{E}_{\mathbf{z}} [\max(0, 1 + D_k(P_k(G(\mathbf{z}))))] \right), \quad (7)$$

$$\mathcal{L}_G = \sum_k -\mathbb{E}_{\mathbf{z}} [D_k(P_k(G(\mathbf{z})))]. \quad (8)$$

During the training process, we apply image blurring through a Gaussian filter with  $\sigma = 2$  pixels for the initial 200k images. Such discriminator blurring (Karras et al., 2021; Sauer et al., 2022) mitigates the discriminator from focusing on high-frequency details in the early stages of training. In addition, following (Sauer et al., 2021), we use spectral normalization without gradient penalties.

While conventional GAN approaches operate in high-dimensional latent spaces, recent studies suggest that natural image datasets possess a relatively lower intrinsic dimensionality (Pope et al., 2021). StyleGAN-XL (Sauer et al., 2022) uses a low dimensional latent space and works more optimally. With these insights, we reduce the dimension of the latent code  $\mathbf{z}$  to 64, resulting in lower FID (Config C).

#### 3.2. Discriminator consistency

In HP-GAN, discriminator consistency is introduced as a regularization technique to minimize the difference in discriminator outputs. The purpose of this consistency loss is to align the outputs of discriminators operating on feature maps extracted from different networks. The discriminator consistency loss is formulated as follows:

$$\mathcal{L}_{DC}(\mathbf{x}) = \mathbb{E}_{\mathbf{x}} \left[ \left( \sum_{k_{\text{CNN}}} D_k(P_k(\mathbf{x})) - \sum_{k_{\text{ViT}}} D_k(P_k(\mathbf{x})) \right)^2 \right], \quad (9)$$

where  $k_{\text{CNN}}$  and  $k_{\text{ViT}}$  denote the set of  $k$  discriminators within the CNN and ViT feature networks, respectively, as illustrated in Fig. 1. During training, this loss is applied on both real images  $\mathbf{x}$  and generated images  $G(\mathbf{z})$ .

Discriminator consistency leverages the distinctive image processing capabilities of CNNs and ViTs. Due to the structure of the networks and the inherent differences in how each network processes the image between CNN and ViT, disparities in discriminator outputs are inevitable. However, by aligning these outputs, we can achieve mutual agreement between discriminators regarding the quality of generated images. Rather than enforcing feature alignment, which may suppress architectural diversity, we regularize the outputs to retain complementary strengths while promoting agreement in judgment.

The important aspect of discriminator consistency is that discriminators that evaluate feature maps from different network architectures can learn from each other. Unlike previous methods (Raghu et al., 2021; Sauer et al., 2022), which train multiple discriminators on separate projected features without explicitly coupling them, our discriminator consistency loss enforces agreement between discriminators. This leads to more stable and robust training by preventing any one discriminator from dominating the feedback process. Encouraging discriminators to converge on a consensus about image quality provides the generator with more uniform and coherent feedback, thereby enhancing the overall stability and performance of the GAN.

To empirically demonstrate the benefits of discriminator consistency, we plot the signed real logits in Fig. 2. The signed real logits, denoted as  $\text{sign}(D(\mathbf{x}))$ , represent the proportion of the training set that receives positive discriminator outputs. This metric serves as a useful heuristic for assessing discriminator overfitting (Karras et al., 2020a;

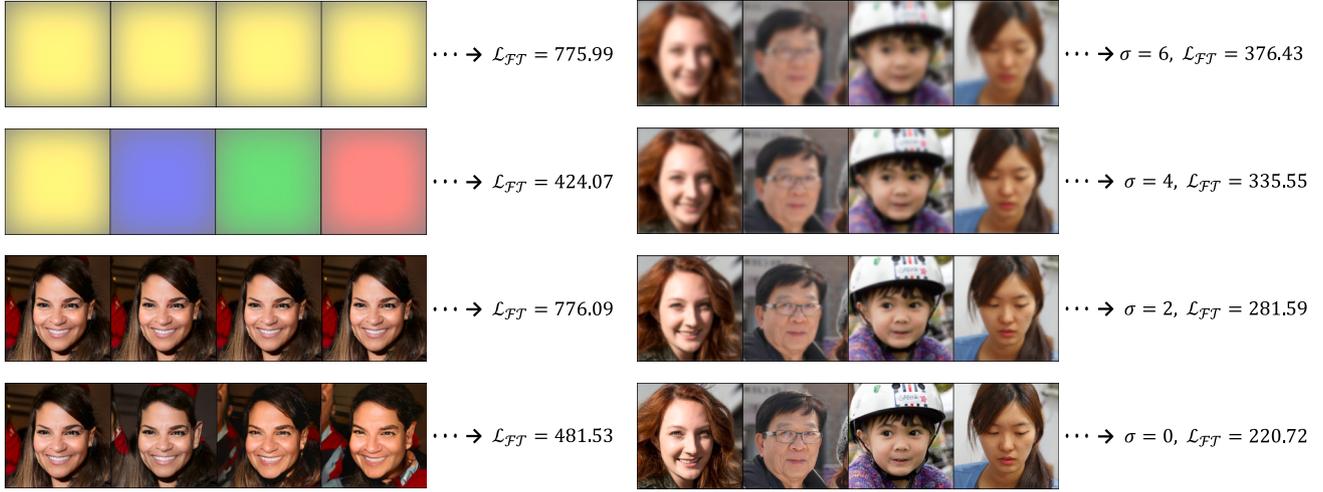


Fig. 4.  $\mathcal{L}_{FT}$  in batches of images with various settings. Rows 1 and 2 in the left column are the results when single color image is repeated, and when multiple color images are in the batch. Rows 3 and 4 in the left column are when single face image is repeated, and when the batch consists of similar but different images generated with noise perturbation. The right column is the comparison according to the degree of Gaussian blurring. We compute 100 times with different random augmentation and report the average for each result.

Sauer et al., 2021; Zhang et al., 2025). A lower  $sign(D(x))$  value, as observed in Config D, indicates reduced overfitting, suggesting that discriminator consistency contributes to a more generalized and stable training process.

In addition, as shown in Table 1, the recall metric, which distinguishes mode-collapse (Borji, 2022), also improves with the application of discriminator consistency. By implementing discriminator consistency, HP-GAN not only stabilizes the training process but also improves the quality of the synthesized images. A more detailed discussion is provided in Section 5.3.

### 3.3. FakeTwins

Rather than imposing additional losses on the discriminator to indirectly improve the performance of the generator, we propose a direct approach that leverages pretrained networks within the GAN training. We introduce FakeTwins, an SSL based mechanism designed to enhance the generative capabilities of the generator.

To effectively utilize pretrained networks, our approach is grounded in the principles of Barlow Twins (Zbontar et al., 2021). The Barlow Twins method promotes feature decorrelation by ensuring that the cross-correlation matrix of the output features approximates the identity matrix. This encourages orthogonality among features, thereby minimizing redundancy and enhancing the diversity of the learned representations.

We hypothesize that obtaining features, normalizing, computing the cross-correlation matrix, and calculating the Barlow Twins loss using a pretrained network would result in lower output loss values for batches containing a diverse set of images. This insight motivates the development of FakeTwins, which applies self-supervised learning (SSL) leveraging pretrained networks to enhance the diversity of the generated images.

In our implementation, we leverage pretrained networks with random projections as encoders for self-supervised learning. Notably, these encoders remain fixed during training, and only the linear head network undergoes training. Representations extracted from two different encoders, subjected to global average pooling, are concatenated to form the final representations that feed into the linear head network.

The FakeTwins loss is specifically applied to perturbed fake images, thereby exclusively targeting the generator. And noise-related latent augmentation (Li et al., 2022) is utilized for noise perturbation, which alleviates the discontinuity of latent space. The formulation of

Table 2

Results on benchmark datasets: FFHQ, LSUN-Bedroom and LSUN-Church. FID is reported as the evaluation metric. The best results are in bold.

Method	FFHQ	Bedroom	Church
ADM (Dhariwal & Nichol, 2021)	2.57	1.90	-
LDM (Rombach et al., 2022)	4.98	2.95	4.02
SAGAN (Zhang et al., 2019)	16.21	14.06	6.15
FastGAN (Liu et al., 2021)	12.69	8.24	8.43
GANformer (Hudson & Zitnick, 2021)	7.42	6.15	5.47
StyleGAN2 (Karras et al., 2019)	4.86	4.01	4.54
StyleGAN3 (Karras et al., 2021)	3.80	-	-
Dual Contrastive (Yu et al., 2021)	4.63	3.31	3.39
Diffusion GAN (Wang et al., 2023b)	3.73	1.43	1.85
Projected GAN (Sauer et al., 2021)	3.39	1.52	1.59
Vision-aided GAN (Sauer et al., 2021)	3.30	-	1.72
DynamicD (Yang et al., 2022)	3.53	4.01	3.87
CLR-GAN (Sun et al., 2024)	3.37	-	3.43
GGDR (Lee et al., 2022)	3.25	3.71	2.81
GLead (Bai et al., 2023)	2.90	2.72	2.15
HP-GAN	<b>1.69</b>	<b>1.19</b>	<b>1.44</b>

the FakeTwins loss is as follows:

$$A' = T_1(G(\mathbf{z})), \quad B' = T_2(G(\mathbf{z} + \hat{\epsilon})), \quad (10)$$

$$V^{A'} = \text{Concat}(P_{\text{CNN}}(A'), P_{\text{ViT}}(A')), \quad (11)$$

$$V^{B'} = \text{Concat}(P_{\text{CNN}}(B'), P_{\text{ViT}}(B')), \quad (12)$$

$$\mathcal{L}_{FT} = \mathcal{L}_{BT}(\phi(V^{A'}), \phi(V^{B'})), \quad (13)$$

where  $\hat{\epsilon} = I_1 \cdot |\mathbf{z}|$  and  $I_1$  is a coefficient set to 0.1.  $T_1$  and  $T_2$  are different augmentations. Here,  $P_{\text{CNN}}$  and  $P_{\text{ViT}}$  represent feature projectors composed of pretrained networks.  $\phi$  denotes the linear head network.

Fig. 3 illustrates the reduction in FID as training progresses, alongside the decreasing FakeTwins loss. Comparing Config D and E, it is evident that FakeTwins promotes faster convergence and improves FID.

The FakeTwins loss according to the composition of the images within a batch is shown in Fig. 4, averaged over 100 random augmentations. We compare  $\mathcal{L}_{FT}$  with respect to the quality, diversity, and degree of Gaussian blurring of the images within the batch. The results indicate that when all images in a batch are identical as shown in the 1st and 3rd rows of the left column, the FakeTwins loss is relatively high, regardless of whether the image is a simple color or a normal human face. Conversely, when the batch comprises varied images, even if they are



Fig. 5. Generated images on benchmark dataset. Categories (top to bottom): FFHQ, LSUN-Bedroom and LSUN-Church.

simple color images (2nd row of the left column), the FakeTwins loss is lower. Additionally, when the batch consists of similar images generated with noise perturbation rather than identical images, the FakeTwins loss is reduced but remains higher than when the images are diverse (4th row). The right column shows that the less the Gaussian blur in images, the lower the FakeTwins loss. This finding suggests that images with more detail result in a lower loss value.

In summary, the diversity of images within the batch, along with the abundance of features and details within individual images, leads to a decrease in FakeTwins loss, thereby improving the generator’s ability to produce diverse and high quality images.

### 3.4. Overall objective function

The overall objective function of HP-GAN is formulated as follows:

$$\mathcal{L}'_D = \mathcal{L}_D + \lambda'_D \mathcal{L}_{DC}(G(\mathbf{z})) + \lambda'_D \mathcal{L}_{DC}(\mathbf{x}), \quad (14)$$

$$\mathcal{L}'_G = \mathcal{L}_G + \lambda_G \mathcal{L}_{DC}(G(\mathbf{z})) + \lambda_f \mathcal{L}_{FT}, \quad (15)$$

where  $\lambda'_D$ ,  $\lambda'_D$ ,  $\lambda_G$  and  $\lambda_f$  are constants that determine the relative weighting of each component. In  $\mathcal{L}'_D$ , the discriminator consistency loss  $\mathcal{L}_{DC}$  is applied both to real images  $\mathcal{L}_{DC}(\mathbf{x})$  and to generated images  $\mathcal{L}_{DC}(G(\mathbf{z}))$ . In  $\mathcal{L}'_G$ ,  $\mathcal{L}_{DC}(G(\mathbf{z}))$  is applied to generated images, since the generator is updated through the consistency of discriminator outputs on generated samples.  $\mathcal{L}_{FT}$  is the FakeTwins loss, which is applied through generated images to the generator.

## 4. Experiments

In this section, we evaluate the performance of our proposed HP-GAN by benchmarking it against state-of-the-art models across multiple datasets. After introducing the implementation details, we present comparisons on various datasets. Furthermore, ablation studies investigate the smoothness of the latent space using Perceptual path length (PPL) and the impact of training with limited data by assessing performance on a subset of the FFHQ dataset. Additionally, we explore the effects of different combinations of feature networks by comparing configurations that use only CNNs with those that combine CNN and ViT.

### 4.1. Experimental setup

**Datasets.** For comparison, all the datasets except AFHQ were resized to  $256^2$ . We evaluate our approach across a wide range of datasets to ensure comprehensive evaluation.

Table 3

Results on few-shot Datasets. FID is reported as the evaluation metric. The best results are in bold.

Method	100-shot			AnimalFace	
	Obama	Grumpy Cat	Panda	Cat	Dog
EDM + DA (Karras et al., 2022)	37.10	29.94	10.81	36.88	57.14
Patch Diffusion (Wang et al., 2023a)	41.47	30.89	13.25	43.71	72.17
StyleGAN2 (Karras et al., 2020b)	80.20	48.90	34.27	71.71	130.2
Scale/shift (Noguchi & Harada, 2019)	50.72	34.20	21.38	54.83	83.04
MineGAN (Wang et al., 2020)	50.63	34.54	14.84	54.45	93.03
TFGAN (Wang et al., 2018b)	48.73	34.06	23.20	52.61	82.38
TFGAN + DA (Zhao et al., 2020b)	39.85	29.77	17.12	49.10	65.57
FreezeD (Mo et al., 2020)	41.87	31.22	17.95	47.70	70.46
StyleGAN2-DA (Zhao et al., 2020b)	46.87	27.08	12.06	42.44	58.85
FastGAN (Liu et al., 2021)	41.05	26.65	10.03	35.11	50.66
StyleGAN2-ADA (Karras et al., 2020a)	45.69	26.62	12.90	40.77	56.83
LeCam GAN (Tseng et al., 2021)	33.16	24.93	10.16	34.18	54.88
GenCo (Cui et al., 2022)	32.21	17.79	9.49	30.89	49.63
InsGen (Yang et al., 2021)	32.42	22.01	9.85	33.01	44.93
FakeCLR (Li et al., 2022)	26.95	19.56	8.42	26.34	42.02
Projected GAN (Sauer et al., 2021)	11.21	15.80	3.98	18.01	17.88
Diffusion GAN (Wang et al., 2023b)	10.54	15.13	3.39	17.86	17.22
DANI (Zhang et al., 2024)	10.08	14.92	3.04	17.72	16.81
QADRS (Zhang et al., 2025)	<b>9.93</b>	14.72	<b>3.08</b>	16.94	16.57
HP-GAN	10.30	<b>13.21</b>	3.34	<b>14.34</b>	<b>15.68</b>



Fig. 6. Generated images on few-shot datasets. Top to bottom (Obama, Grumpy Cat, Panda, AnimalFace Cat, AnimalFace Dog).

- Large datasets: We present the comparisons on large datasets such as FFHQ (70k images) (Karras et al., 2019), LSUN-Bedroom (3M images) (Yu et al., 2015), LSUN-Church (126k images) (Yu et al., 2015), CLEVR (70k images) (Johnson et al., 2017), and Cityscapes (25k images) (Cordts et al., 2016), with a primary focus on FFHQ, LSUN-Bedroom, and LSUN-Church as benchmark datasets.
- Small datasets: To evaluate the performance on small datasets, datasets used for the evaluation include art paintings from WikiArt

**Table 4**

Results on AFHQ. FID is reported as the evaluation metric. The best results are in bold.

Method	AFHQ		
	Cat	Dog	Wild
DDMI (Park et al., 2024)	4.27	8.54	-
StyleGAN2 (Karras et al., 2020b)	5.13	19.4	3.48
ContraD (Jeong & Shin, 2021)	3.82	7.16	2.54
StyleGAN2-ADA (Karras et al., 2020a)	3.55	7.41	3.05
InsGen (Yang et al., 2021)	2.60	5.44	1.77
Vision-aided GAN (Kumari et al., 2022)	2.44	4.60	2.25
Diffusion GAN (Wang et al., 2023b)	2.40	4.83	1.51
Projected GAN (Sauer et al., 2021)	2.16	4.52	2.17
HP-GAN	<b>1.81</b>	<b>3.63</b>	<b>1.18</b>

(1000 images; wikiart.org), Oxford Flowers (1360 images) (Nilsback & Zisserman, 2008), landscape photographs (4319 images; flickr.com), AnimalFace Dog (389 images) (Si & Zhu, 2012), and Pokemon (833 images; pokemon.com).

- Few-shot datasets: Our experiments also explore few-shot settings using datasets as Obama (100 images) (Zhao et al., 2020b), Grumpy Cat (100 images) (Zhao et al., 2020b), Panda (100 images) (Zhao et al., 2020b), AnimalFace Cat (160 images) (Si & Zhu, 2012), and AnimalFace Dog (389 images) (Si & Zhu, 2012).
- AFHQ datasets: We evaluate the performance on the AFHQ dataset (Choi et al., 2020), which consists of around 5000 images per category for cat, dog, and wildlife, all at a resolution of  $512^2$ .

**Implementation.** For the generator  $G$ , we adopted the architecture of FastGAN (Liu et al., 2021; Sauer et al., 2021). And for the discriminator  $D$ , we used the architecture of StyleGAN-XL (Sauer et al., 2022). We used the combination of EfficientNet-lite0 (Tan & Le, 2019) and DeiT-B (Touvron et al., 2021) for the feature networks. The linear head network for FakeTwins had three linear layers, each with 512 output units, and the first two layers were followed by batch normalization and rectified linear units. Exponential moving average of generator weights and Adam (Kingma & Ba, 2015) optimizer were adopted. We employed differentiable data augmentation (Zhao et al., 2020b) to effectively train GAN and also applied it to generate distorted images for FakeTwins. For all datasets, we applied data amplification via x-flips. All models were trained with 20M images to ensure convergence, except for the few-shot datasets, where training was extended to 100M images. The hyperparameters in HP-GAN are set as  $\lambda_D^f = 1$ ,  $\lambda_D^r = 1$ ,  $\lambda_G = 1$  and  $\lambda_f = 0.02$ . All experiments were conducted on 8 NVIDIA RTX A5000 GPUs with 64 batch size.

**Evaluation.** We employed the Fréchet Inception Distance (FID) (Heusel et al., 2017) as the main quantitative assessment metric, because FID tends to correlate with human judgement regarding the quality of synthesized images. Regardless of the number of training data, we computed the FID between 50k generated images and all training images (Karras et al., 2019, 2020b; Liu et al., 2021). Our implementation is based on the official codebase of ADA (Karras et al., 2020a). And we report the lowest (best) FID, as in the notable studies (Sauer et al., 2021, 2022).

Additionally, we compared the precision and recall (Kynkäänniemi et al., 2019) to provide further insights into the fidelity and diversity of the generated images, and Kernel Inception Distance (KID) (Bińkowski et al., 2018), which is more suitable for small datasets.

## 4.2. Results

**Benchmark datasets.** Table 2 presents the comparisons of HP-GAN against various generative models on benchmark datasets. HP-GAN consistently achieves superior performance across all datasets. We compare against both GAN and diffusion methods. Among GANs, we in-



**Fig. 7.** Generated images on AFHQ dataset. Categories (top to bottom): Cat, Dog and Wild.

clude StyleGAN2 (Karras et al., 2019), Projected GAN (Sauer et al., 2021), Diffusion GAN (Wang et al., 2023b), GGDR (Lee et al., 2022), and GLeAD (Bai et al., 2023). For diffusion models, we evaluate against ADM (Dhariwal & Nichol, 2021) and LDM (Rombach et al., 2022). Notably, HP-GAN achieves the lowest FID scores across all datasets, with a particularly significant improvement on FFHQ, reducing the FID to 1.69. Qualitative results are shown in Fig. 5.

**Few-shot datasets.** Few-shot datasets pose significant challenges due to the limited training samples and high risk of overfitting. As shown in Table 3, HP-GAN demonstrates consistently strong performance across all datasets, achieving competitive FID scores. While it does not outperform all methods on every dataset, it achieves state-of-the-art results on AnimalFace Cat and Dog, as well as on Grumpy Cat. Qualitative results for few-shot datasets are provided in Fig. 6, showing that HP-GAN produces visually consistent and diverse generations even under extreme data constraints.

**AFHQ datasets.** At a resolution of  $512^2$ , HP-GAN achieves superior performance across all categories of the AFHQ datasets. It achieves an FID of 1.81 for the Cat category and 3.63 for the Dog category, outperforming all compared methods. Furthermore, HP-GAN attains an FID of 1.18 on the Wild category, highlighting its strong generalization across diverse animal domains. Quantitative results are presented in Table 4 and generated images are shown in Fig. 7.

**Various Datasets.** To assess the robustness of HP-GAN across diverse conditions and dataset sizes, we compare it against strong baselines, including ADA (Karras et al., 2020a), FastGAN (Liu et al., 2021), and Projected GAN (Sauer et al., 2021). The results, evaluated using FID, KID, precision (Pr), and recall (Rec), are summarized in Table 5. For KID, values are scaled by  $10^3$ , and arrows ( $\uparrow/\downarrow$ ) indicate whether higher or lower is better. Representative qualitative results are shown in Fig. 5.

To complement FID, we report Kernel Inception Distance (KID) (Bińkowski et al., 2018), which is an unbiased alternative that tends to be more stable, especially on smaller datasets. Except for two datasets, the KID rankings are consistent with those of FID, indicating that HP-GAN’s performances is not overly tuned to a specific metric.

We further evaluate generation quality and diversity using precision, which measures the fidelity of generated images, and recall, which quantifies the coverage of the real distribution (Kynkäänniemi et al., 2019). These two metrics reflect the trade-off: low precision suggests degraded image quality, while low recall indicates mode collapse (Borji, 2022).

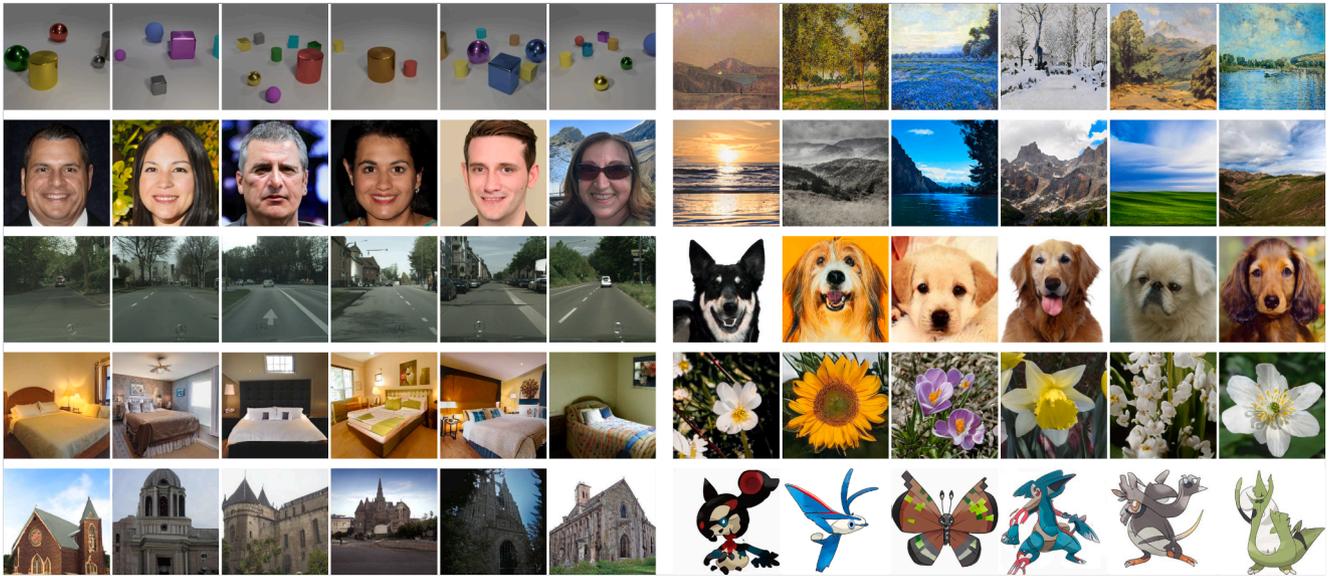


Fig. 8. Generated images with various datasets. *Left*: Large datasets (CLEVR, FFHQ, Cityscapes, Bedroom, Church). *Right*: Small datasets (Art Painting, Landscape, AnimalFace Dog, Flowers, Pokemon).

Some methods achieve better performance in precision compared to HP-GAN. Methods like ADA (Karras et al., 2020a) and FastGAN (Liu et al., 2021), though strong in precision (fidelity), struggle with recall, suggesting susceptibility to mode collapse. In contrast, HP-GAN achieves consistently high recall across most datasets, which explains its strong performance in FID and KID. These results highlight HP-GAN’s ability to maintain high fidelity while preserving sample diversity.

On high-resolution datasets such as AFHQ, both HP-GAN and Projected GAN (Sauer et al., 2021) show slightly reduced precision. However, HP-GAN surpasses Projected GAN in both precision and recall, demonstrating improvements in image quality and diversity as well as further alleviation of mode collapse.

## 5. Ablation study

### 5.1. Perceptual path length (PPL)

PPL (Karras et al., 2019) is a metric introduced to evaluate the smoothness of the latent space, calculating the average LPIPS distances (Zhang et al., 2018b) between generated images when small perturbations is applied to the latent space. This metric also correlates with the consistency and stability of shapes and higher overall image quality (Karras et al., 2020b). As shown in Table 1, we compute PPL in latent space,  $\mathcal{Z}$ , for the entire images without the central crop. Assessments are conducted on configurations where  $\mathbf{z} \in \mathbb{R}^{64}$ .

Our proposed methods, FakeTwins and discriminator consistency, contribute to enhancing the diversity of the generated images. And the improvement of PPL demonstrates that the incorporation of these techniques not only enables the generation of diverse images but also ensures high quality. This indicates that HP-GAN effectively balances the generation of diverse images while maintaining superior image quality.

### 5.2. Data efficiency evaluation

To evaluate the model’s efficiency with limited training data, we conduct experiments using different subset sizes of the FFHQ dataset by randomly sampling with 100, 1k, 2k, and 5k images. Regardless of the number of training data, FID is calculated between 50k generated images and all FFHQ (70k images). The results are shown in Table 6.

HP-GAN consistently outperforms existing methods across all subset sizes, except for the 100 sample case, where it ranks second. The im-

provement in FID is remarkable, demonstrating the data efficiency and effectiveness of the model even when limited data is available. Notably, HP-GAN achieves comparable performance compared to other methods trained on the entire dataset (Table 2), even when only using 2k or 5k images. This highlights the robustness and capability of HP-GAN in achieving high-quality image generation with limited data.

### 5.3. Feature networks combination

The combination of feature networks is an important part of HP-GAN. Our ablation studies on FFHQ dataset, detailed in Table 7, compare two distinct combinations: dual CNNs using EfficientNet-lite0 and ResNet50, and the hybrid of CNN & ViT using EfficientNet-lite0 and DeiT-B.

The most significant improvement is observed when discriminator consistency is incorporated, as seen in Config D and E. The enhanced performance achieved with the CNN and ViT hybrid indicates that discriminators are more effectively trained when leveraging feature networks with different architectures. In contrast, with dual CNNs, the effect of discriminator consistency is small because both networks process images with similar mechanisms.

Furthermore, SSL using both the ViT and CNN simultaneously shows higher performance than using the same network architecture (Song et al., 2023). Additionally, when training networks with distillation, performance is better when the teacher and student networks have different architectures (Touvron et al., 2021). These findings support the advantage of hybrid feature networks, which benefit from the strengths of both CNN and ViT through the application of discriminator consistency, as discussed in Section 3.2.

### 5.4. Choice of SSL objective for FakeTwins

FakeTwins can be implemented with different SSL objectives. We compare several choices on FFHQ-1k, keeping all GAN hyperparameters fixed and changing only the SSL loss applied to the generator. The results are reported in Table 8.

Both VICReg (Bardes et al., 2022) and Barlow Twins (Zbontar et al., 2021) improve over the no SSL baseline, while SimCLR (Chen et al., 2020a) does not. In our framework, the SSL loss is applied directly to the generator, and SimCLR, which explicitly separates samples via negative pairs, may emphasize instance-level discrimination that is less useful

**Table 5**

Comparison experiment results for FID, KID, Precision (Pr), and Recall (Rec) on various datasets. KID indicates the value multiplied by  $10^3$  ( $KID \times 10^3$ ). Best results are in bold, second best are underlined.

Dataset	Method	FID ↓	KID ↓	Pr ↑	Rec ↑
CLEVR	ADA (Karras et al., 2020a)	10.17	8.15	0.373	0.569
	FastGAN (Liu et al., 2021)	3.24	2.64	0.600	0.650
	Projected GAN (Sauer et al., 2021)	<u>0.89</u>	<u>0.51</u>	<u>0.640</u>	<u>0.735</u>
	HP-GAN	<b>0.81</b>	<b>0.42</b>	<b>0.693</b>	<b>0.718</b>
FFHQ	ADA (Karras et al., 2020a)	7.32	1.49	<u>0.669</u>	0.445
	FastGAN (Liu et al., 2021)	12.69	5.34	<b>0.716</b>	0.184
	Projected GAN (Sauer et al., 2021)	<u>3.08</u>	<u>0.44</u>	0.654	<u>0.464</u>
	HP-GAN	<b>1.69</b>	<b>0.19</b>	0.642	<b>0.545</b>
Cityscapes	ADA (Karras et al., 2020a)	8.35	3.34	<b>0.649</b>	0.146
	FastGAN (Liu et al., 2021)	8.78	5.45	0.557	0.227
	Projected GAN (Sauer et al., 2021)	<u>3.41</u>	<u>0.91</u>	<u>0.619</u>	<u>0.361</u>
	HP-GAN	<b>2.11</b>	<b>0.72</b>	0.590	<b>0.469</b>
Bedroom	ADA (Karras et al., 2020a)	11.53	7.42	0.429	0.202
	FastGAN (Liu et al., 2021)	8.24	5.90	0.602	0.189
	Projected GAN (Sauer et al., 2021)	<u>1.52</u>	<u>0.36</u>	<u>0.614</u>	<u>0.346</u>
	HP-GAN	<b>1.19</b>	<u>0.39</u>	<u>0.604</u>	<u>0.321</u>
Church	ADA (Karras et al., 2020a)	5.85	4.70	0.565	0.416
	FastGAN (Liu et al., 2021)	8.43	4.61	<b>0.645</b>	0.207
	Projected GAN (Sauer et al., 2021)	<u>1.59</u>	<u>0.50</u>	<u>0.612</u>	<u>0.438</u>
	HP-GAN	<b>1.44</b>	<b>0.48</b>	0.577	<b>0.459</b>
Art Painting	ADA (Karras et al., 2020a)	43.07	10.23	0.691	<u>0.218</u>
	FastGAN (Liu et al., 2021)	44.02	13.00	<b>0.858</b>	0.044
	Projected GAN (Sauer et al., 2021)	<u>27.96</u>	<u>1.25</u>	0.762	<b>0.239</b>
	HP-GAN	<b>26.31</b>	<b>0.62</b>	<u>0.784</u>	0.131
Landscape	ADA (Karras et al., 2020a)	15.99	4.39	0.709	0.213
	FastGAN (Liu et al., 2021)	16.44	3.40	<u>0.768</u>	0.160
	Projected GAN (Sauer et al., 2021)	<u>6.92</u>	<u>1.30</u>	<u>0.774</u>	<u>0.258</u>
	HP-GAN	<b>6.08</b>	<b>0.28</b>	0.716	<b>0.425</b>
AnimalFace Dog	ADA (Karras et al., 2020a)	60.90	22.52	0.841	0.036
	FastGAN (Liu et al., 2021)	62.11	22.11	0.849	0.015
	Projected GAN (Sauer et al., 2021)	<u>17.88</u>	<u>0.03</u>	<u>0.998</u>	<u>0.095</u>
	HP-GAN	<b>16.38</b>	<u>0.25</u>	<u>0.993</u>	<b>0.242</b>
Flowers	ADA (Karras et al., 2020a)	21.66	3.56	0.731	0.095
	FastGAN (Liu et al., 2021)	26.23	6.61	0.611	<u>0.100</u>
	Projected GAN (Sauer et al., 2021)	<u>13.86</u>	<u>0.38</u>	<u>0.816</u>	0.058
	HP-GAN	<b>12.03</b>	<b>0.30</b>	<b>0.818</b>	<b>0.107</b>
Pokemon	ADA (Karras et al., 2020a)	40.38	13.49	0.735	0.197
	FastGAN (Liu et al., 2021)	81.86	80.30	0.731	0.004
	Projected GAN (Sauer et al., 2021)	<u>26.36</u>	<u>1.32</u>	<b>0.809</b>	<u>0.259</u>
	HP-GAN	<b>23.62</b>	<b>1.12</b>	<u>0.768</u>	<b>0.310</b>
AFHQ Cat	ADA (Karras et al., 2020a)	3.55	0.63	<u>0.767</u>	0.411
	FastGAN (Liu et al., 2021)	4.69	1.72	<b>0.784</b>	0.305
	Projected GAN (Sauer et al., 2021)	<u>2.16</u>	<u>0.16</u>	0.693	<u>0.565</u>
	HP-GAN	<b>1.81</b>	<b>0.08</b>	0.741	<b>0.580</b>
AFHQ Dog	ADA (Karras et al., 2020a)	7.40	1.21	<b>0.753</b>	0.470
	FastGAN (Liu et al., 2021)	13.09	5.51	<u>0.746</u>	0.380
	Projected GAN (Sauer et al., 2021)	<u>4.52</u>	<u>0.80</u>	<u>0.718</u>	<u>0.643</u>
	HP-GAN	<b>3.63</b>	<b>0.37</b>	0.737	<b>0.677</b>
AFHQ Wild	ADA (Karras et al., 2020a)	3.05	0.47	<b>0.765</b>	0.137
	FastGAN (Liu et al., 2021)	3.14	0.74	<u>0.761</u>	0.201
	Projected GAN (Sauer et al., 2021)	<u>2.17</u>	<u>0.12</u>	0.705	<u>0.292</u>
	HP-GAN	<b>1.18</b>	<b>0.02</b>	0.737	<b>0.367</b>

from the generator’s perspective and is therefore not necessarily well suited.

In contrast, VICReg and Barlow Twins are negative-free, information-maximization methods that encourage invariance between augmented views while increasing the information content of the embeddings. This type of objective is more compatible with FakeTwins, as it promotes informative and diverse feature representations of generated images. Consequently, we adopt Barlow Twins as the SSL objective for FakeTwins.

**Table 6**

Ablation study on FFHQ Subsets. FID is reported as the evaluation metric. The best results are in bold.

Method	FFHQ			
	100	1k	2k	5k
EDM + DA (Karras et al., 2022)	50.73	30.75	27.17	24.51
Patch Diffusion (Wang et al., 2023a)	44.45	28.03	25.32	22.63
StyleGAN2 (Karras et al., 2020b)	179.00	100.16	54.30	49.68
GenCo (Cui et al., 2022)	148.00	65.31	47.32	27.96
StyleGAN2-DA (Zhao et al., 2020b)	61.91	25.66	24.43	10.45
StyleGAN2-ADA (Karras et al., 2020a)	85.80	21.29	15.39	10.96
APA (Jiang et al., 2021)	65.00	18.89	16.90	8.38
InsGen (Yang et al., 2021)	45.75	18.21	11.47	7.83
FakeCLR (Li et al., 2022)	42.56	15.92	9.90	7.25
DANI (Zhang et al., 2024)	23.98	10.81	7.73	6.20
Projected GAN (Sauer et al., 2021)	26.25	11.12	8.25	6.85
Diffusion GAN (Wang et al., 2023b)	25.47	8.76	7.99	6.59
QADDRS (Zhang et al., 2025)	<b>19.89</b>	6.57	5.18	4.69
HP-GAN	22.20	<b>5.22</b>	<b>3.62</b>	<b>2.86</b>

**Table 7**

Ablation study for feature networks on FFHQ dataset. The combination of CNN & CNN and CNN & ViT. FID is reported as the evaluation metric. The best results are in bold.

Method	Feature Networks	
	CNN & CNN	CNN & ViT
StyleGAN-XL (Sauer et al., 2022)	-	2.19
Config C	2.66	2.29
Config D	2.61	1.86
Config E	2.25	<b>1.69</b>

**Table 8**

Ablation study on FFHQ-1k with different SSL objectives used in FakeTwins. The best results are in bold.

Method	FID ↓
Config D (no SSL loss)	5.65
FakeTwins w/ SimCLR (Chen et al., 2020a)	5.83
FakeTwins w/ VICReg (Bardes et al., 2022)	5.58
FakeTwins w/ Barlow Twins (Zbontar et al., 2021)	<b>5.22</b>

## 6. Conclusion

In this study, we introduce HP-GAN, a novel approach to GANs that effectively harnesses the power of pretrained networks and self-supervised learning to enhance image synthesis. Our method leverages FakeTwins to train generator, thereby improving the diversity and quality of images. Furthermore, we introduce discriminator consistency, which aligns the outputs of multiple discriminators with different feature network architectures, promoting more stable and robust training. Our extensive experiments across a variety of datasets, including large, small, and few-shot settings, demonstrate that HP-GAN consistently outperforms state-of-the-art models. We show that HP-GAN maintains high performance even with limited training data, emphasizing its efficiency and robustness. The ablation studies provide evidence that the hybrid use of CNN and ViT in our feature networks significantly enhances performance, validating the benefits of our approach. Overall, HP-GAN advances the field of generative modeling by integrating pretrained networks and innovative self-supervised learning strategies, offering a powerful tool for generating high-quality images across different contexts. This work sets a new precedent for future exploration and development of leveraging pretrained networks.

## CRedit authorship contribution statement

**Geonhui Son:** Writing - review & editing, Writing - original draft, Visualization, Validation, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Jeong Ryong Lee:** Conceptualization; **Dosik Hwang:** Project administration, Funding acquisition, Conceptualization.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

This work was supported in part by the [National Research Foundation of Korea](#) (NRF) grant funded by the Korea government(MSIT) (RS-2025-02215070, RS-2025-02217919, RS-2025-16070382); in part by the MSIT(Ministry of Science, ICT), Korea, under the Global Research Support Program in the Digital Field program (RS-2024-00436680) supervised by the IITP ( & Communications Technology Planning & Evaluation); in part by Artificial Intelligence Graduate School Program at Yonsei University (RS-2020-II201361); in part by the KIST Institutional Program (2E33801, 2E33800); in part by the Yonsei Signature Research Cluster Program of 2024 (2024-22-0161).

## Appendix A. Fine-tuning from Natural Images to Medical Images.

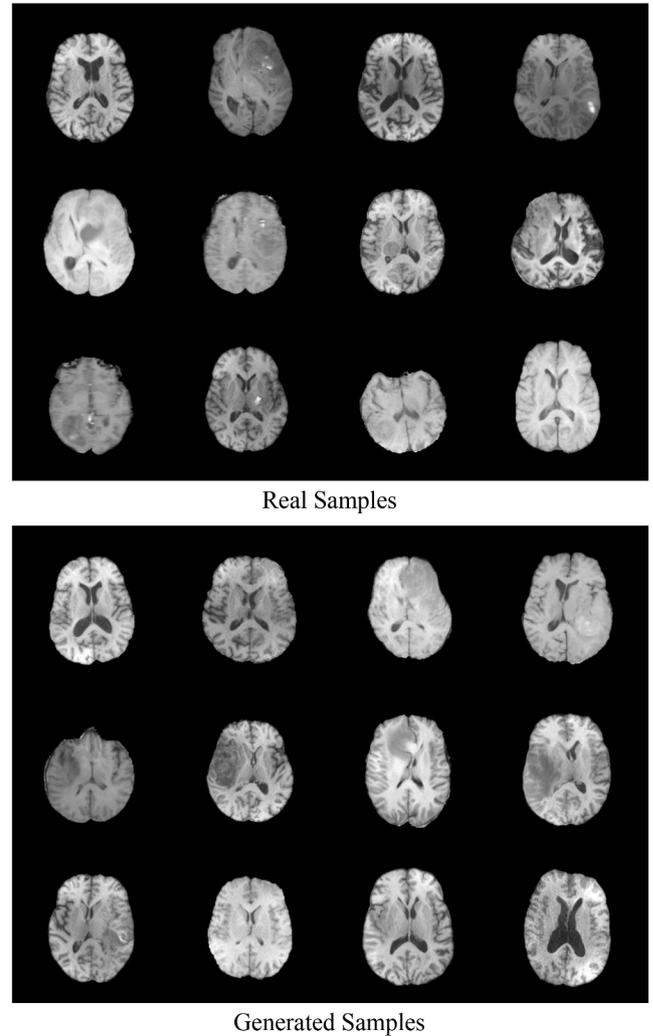
Generative models for medical imaging, including GAN- and diffusion-based approaches, have attracted considerable attention for data augmentation, disease diagnosis, anatomical segmentation, and anomaly detection (Sharma & Hamarneh, 2019; Shin et al., 2025a,b).

We investigate whether HP-GAN can be transferred from natural-image pretraining to medical imagery and whether domain mismatch in the pretrained feature networks is a problem. To this end, we conduct an experiment on the BraTS2021 (Baid et al., 2021) dataset using only the T1 contrast. We extract 1,151 training slices and compare two settings: (i) training HP-GAN from scratch on BraTS, and (ii) fine-tuning HP-GAN initialized from our FFHQ-pretrained model (natural face images).

As shown in Table A.9, fine-tuning from FFHQ leads to an improvement in FID compared to training from scratch, indicating that HP-GAN benefits from natural-image pretraining even when the target domain is medical imagery. Fig. A.9 provides qualitative examples for the fine-tuned model.

**Table A.9**  
Quantitative results on BraTS2021 T1 for HP-GAN trained from scratch and fine-tuned from an FFHQ-pretrained model.

HP-GAN setting	FID ↓
Trained from scratch	27.73
Fine-tuned from FFHQ	<b>9.39</b>



**Fig. A.9.** Qualitative results on BraTS2021 T1 contrast. Top: real training samples. Bottom: images generated by HP-GAN fine-tuned from FFHQ, demonstrating successful transfer from natural to medical images.

## Appendix B. Alternative Feature Networks

In the main experiments, we adopt EfficientNet-lite0 and DeiT-B, which were trained for classification, as the default feature networks following StyleGAN-XL (Sauer et al., 2022). To examine the dependence of HP-GAN on the specific choice of feature networks, we evaluate several alternative combinations on FFHQ, varying the pretraining strategy (supervised classification, self-supervised DINO (Caron et al., 2021), and CLIP (Radford et al., 2021)). The configurations and FID scores are summarized in Table B.10.

All these configurations yield strong performance, with relatively small differences in FID. This suggests that HP-GAN is not tied to a particular feature-network architecture or pretraining objective: both supervised and self-supervised ViT backbones (e.g., DINO) and CLIP-pretrained CNNs work well, as long as they provide sufficiently informative features. This observation is consistent with prior findings in

**Table B.10**

Quantitative results on FFHQ for HP-GAN with different combinations of feature networks and pretraining objectives. The feature networks are trained with supervised classification (CLS), self-supervised DINO, and CLIP pretraining.

CNN	ViT	FID ↓
ResNet-50 (CLIP)	ViT-B (DINO)	2.03
ResNet-50 (CLIP)	DeiT-B (CLS)	1.91
EfficientNet-lite0 (CLS)	ViT-B (DINO)	1.88
EfficientNet-lite0 (CLS)	DeiT-B (CLS)	<b>1.69</b>

Projected GAN (Sauer et al., 2021) and StyleGAN-XL (Sauer et al., 2022), where pretrained feature networks perform well regardless of training data, pretraining objective, or network architecture.

## References

- Arjovsky, M., & Bottou, L. (2017). Towards principled methods for training generative adversarial networks. In *International conference on learning representations*.
- Arjovsky, M., Chintala, S., & Bottou, L. (2017). Wasserstein generative adversarial networks. In *International conference on machine learning* (pp. 214–223).
- Bachman, P., Hjelm, R. D., & Buchwalter, W. (2019). Learning representations by maximizing mutual information across views. In *Advances in neural information processing systems* (pp. 15509–15519).
- Bai, Q., Yang, C., Xu, Y., Liu, X., Yang, Y., & Shen, Y. (2023). Glead: Improving gans with a generator-leading task. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 12094–12104).
- Baid, U., Ghodasara, S., Mohan, S., Bilello, M., Calabrese, E., Colak, E., Farahani, K., Kalpathy-Cramer, J., Kitamura, F. C., Pati, S. et al. (2021). The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. arXiv preprint arXiv:2107.02314, .
- Bardes, A., Ponce, J., & LeCun, Y. (2022). Vicreg: Variance-invariance-covariance regularization for self-supervised learning. In *International conference on learning representations*.
- Bellemare, M. G., Danihelka, I., Dabney, W., Mohamed, S., Lakshminarayanan, B., Hoyer, S., & Munos, R. (2017). The cramer distance as a solution to biased wasserstein gradients. arXiv:1705.10743, .
- Bińkowski, M., Sutherland, D. J., Arbel, M., & Gretton, A. (2018). Demystifying mmd gans. In *International conference on learning representations*.
- Borji, A. (2022). Pros and cons of GAN evaluation measures: New developments. *Computer Vision and Image Understanding*, 215, 103329.
- Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., & Joulin, A. (2020). Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, (pp. 9912–9924).
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., & Joulin, A. (2021). Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 9630–9640).
- Casanova, A., Careil, M., Verbeek, J., Drozdal, M., & Romero Soriano, A. (2021). Instance-conditioned gan. *Advances in Neural Information Processing Systems*, (pp. 27517–27529).
- Chen, Q., & Koltun, V. (2017). Photographic image synthesis with cascaded refinement networks. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 1511–1520).
- Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. E. (2020a). A simple framework for contrastive learning of visual representations. In *International conference on machine learning* (pp. 1597–1607).
- Chen, X., Fan, H., Girshick, R., & He, K. (2020b). Improved baselines with momentum contrastive learning. arXiv:2003.04297, .
- Chen, X., & He, K. (2021). Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 15750–15758).
- Choi, Y., Uh, Y., Yoo, J., & Ha, J. (2020). StarGAN v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 8188–8197).
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T.,ENZWEILER, M., Benenson, R., Franke, U., Roth, S., & Schiele, B. (2016). The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 3213–3223).
- Cubuk, E. D., Zoph, B., Mane, D., Vasudevan, V., & Le, Q. V. (2018). Autoaugment: Learning augmentation policies from data. arXiv:1805.09501, .
- Cubuk, E. D., Zoph, B., Shlens, J., & Le, Q. V. (2020). Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshop* (pp. 702–703).
- Cui, K., Huang, J., Luo, Z., Zhang, G., Zhan, F., & Lu, S. (2022). Genco: Generative co-training for generative adversarial networks with limited data. In *Proceedings of the AAAI conference on artificial intelligence* (pp. 499–507).
- van den, O. A., Li, Y., & Vinyals, O. (2018). Representation learning with contrastive predictive coding. arXiv:1807.03748, .
- Deshpande, I., Zhang, Z., & Schwing, A. G. (2018). Generative modeling using the sliced wasserstein distance. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 3483–3491).
- Dhariwal, P., & Nichol, A. Q. (2021). Diffusion models beat GANs on image synthesis. In *Advances in neural information processing systems* (pp. 8780–8794).
- Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., & Darrell, T. (2014). DeCAF: A deep convolutional activation feature for generic visual recognition. In *International conference on machine learning* (pp. 647–655).
- Dosovitskiy, A., & Brox, T. (2016). Generating images with perceptual similarity metrics based on deep networks. In *Advances in neural information processing systems* (pp. 658–666).
- Ermolov, A., Siarohin, A., Sangineto, E., & Sebe, N. (2021). Whitening for self-supervised representation learning. In *International conference on machine learning* (pp. 3015–3024). (vol. 139).
- Fedus, W., Rosca, M., Lakshminarayanan, B., Dai, A. M., Mohamed, S., & Goodfellow, I. J. (2018). Many paths to equilibrium: GANs do not need to decrease a divergence at every step. In *International conference on learning representations*.
- Fu, S., Tamir, N., Sundaram, S., Chai, L., Zhang, R., Dekel, T., & Isola, P. (2023). Dream-sim: Learning new dimensions of human visual similarity using synthetic data. arXiv preprint arXiv:2306.09344, .
- Gatys, L. A., Ecker, A. S., & Bethge, M. (2016). Image style transfer using convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 2414–2423).
- Gidaris, S., Bursuc, A., Puy, G., Komodakis, N., Cord, M., & Pérez, P. (2021). Obow: On-line bag-of-visual-words generation for self-supervised learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 6830–6840).
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. C., & Bengio, Y. (2020). Generative adversarial networks. *Communications of the ACM*, 63(11), 139–144.
- Gou, Y., Li, M., Zhang, Y., He, Z., & He, Y. (2024). Few-shot image generation with reverse contrastive learning. *Neural Networks*, 169, 154–164.
- Grill, J., Strub, F., Althé, F., Tallec, C., Richemond, P. H., Buchatskaya, E., Doersch, C., Pires, B. Á., Guo, Z., Azar, M. G., Piot, B., Kavukcuoglu, K., Munos, R., & Valko, M. (2020). Bootstrap your own latent - a new approach to self-supervised learning. In *Advances in neural information processing systems* (pp. 21271–21284).
- Guo, S., Liu, Z., Yang, Z., Lee, C. H., Lv, Q., & Shen, L. (2025). Multi-scale multi-object semi-supervised consistency learning for ultrasound image segmentation. *Neural Networks*, 184, 107095.
- Ham, H., Jun, T. J., & Kim, D. (2020). Unbalanced GANs: Pre-training the generator of generative adversarial network using variational autoencoder. arXiv:2002.02112, .
- He, K., Fan, H., Wu, Y., Xie, S., & Girshick, R. (2020). Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 9729–9738).
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 770–778).
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., & Hochreiter, S. (2017). Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems* (pp. 6626–6637).
- Hou, L., Shen, H., Cao, Q., & Cheng, X. (2021). Self-supervised gans with label augmentation. *Advances in Neural Information Processing Systems*, 34, 13019–13031.
- Hudson, D. A., & Zitnick, L. (2021). Generative adversarial transformers. In *International conference on machine learning* (pp. 4487–4499).
- Huh, M., Agrawal, P., & Efros, A. A. (2016). What makes imagenet good for transfer learning? arXiv:1608.08614, .
- Jeong, J., & Shin, J. (2021). Training GANs with stronger augmentations via contrastive discriminator. In *International conference on learning representations*.
- Jiang, L., Dai, B., Wu, W., & Loy, C. C. (2021). Deceive d: adaptive pseudo augmentation for GAN training with limited data. In *Advances in neural information processing systems* (pp. 21655–21667).
- Johnson, J., Alahi, A., & Fei-Fei, L. (2016). Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision* (pp. 694–711).
- Johnson, J., Hariharan, B., van der Maaten, L., Fei-Fei, L., Zitnick, C. L., & Girshick, R. B. (2017). CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 2901–2910).
- Kang, M., & Park, J. (2020). Contragan: Contrastive learning for conditional image generation. *Advances in Neural Information Processing Systems*, (pp. 21357–21369).
- Kang, M., Shim, W., Cho, M., & Park, J. (2021). Rebooting acgan: Auxiliary classifier gans with stable training. *Advances in Neural Information Processing Systems*, (pp. 23505–23518).
- Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2018). Progressive growing of GANs for improved quality, stability, and variation. In *International conference on learning representations*.
- Karras, T., Aittala, M., Aila, T., & Laine, S. (2022). Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35, 26565–26577.
- Karras, T., Aittala, M., Hellsten, J., Laine, S., Lehtinen, J., & Aila, T. (2020a). Training generative adversarial networks with limited data. In *Advances in neural information processing systems* (pp. 12104–12114).
- Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., & Aila, T. (2021). Alias-free generative adversarial networks. In *Advances in neural information processing systems* (pp. 852–863).
- Karras, T., Laine, S., & Aila, T. (2019). A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 4401–4410).

- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., & Aila, T. (2020b). Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 8110–8119).
- Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., & Krishnan, D. (2020). Supervised contrastive learning. In *Advances in neural information processing systems* (pp. 18661–18673).
- Kim, H., Shin, Y., & Hwang, D. (2023). Dimix: Disentangle-and-mix based domain generalizable medical image segmentation. In *International conference on medical image computing and computer-assisted intervention* (pp. 242–251). Springer.
- Kim, S., Jang, H., Hong, S., Hong, Y. S., Bae, W. C., Kim, S., & Hwang, D. (2021). Fat-saturated image generation from multi-contrast MRIs using generative adversarial networks with bloch equation-based autoencoder regularization. *Medical Image Analysis*, 73, 102198.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *International conference on learning representations*.
- Kornblith, S., Shlens, J., & Le, Q. V. (2019). Do better imagenet models transfer better? In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 2661–2671).
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2017). Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6), 84–90.
- Kumari, N., Zhang, R., Shechtman, E., & Zhu, J.-Y. (2022). Ensembling off-the-shelf models for GAN training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 10651–10662).
- Kunkel, L., & Trabs, M. (2025). A wasserstein perspective of vanilla GANs. *Neural Networks*, 181, 106770.
- Kynkäänniemi, T., Karras, T., Laine, S., Lehtinen, J., & Aila, T. (2019). Improved precision and recall metric for assessing generative models. *Advances in Neural Information Processing Systems*, 32.
- Ledig, C., Theis, L., Huszar, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A. P., Tejani, A., Totz, J., Wang, Z., & Shi, W. (2017). Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 4681–4690).
- Lee, G., Kim, H., Kim, J., Kim, S., Ha, J.-W., & Choi, Y. (2022). Generator knows what discriminator should learn in unconditional gans. In *European conference on computer vision* (pp. 406–422). Springer.
- Lee, J. R., Son, G., & Hwang, D. (2025). Fedi: Feature disentanglement for self-supervised learning. *Pattern Recognition*, (p. 112619).
- Li, C., Chang, W., Cheng, Y., Yang, Y., & Póczos, B. (2017). MMD GAN: Towards deeper understanding of moment matching network. In *Advances in neural information processing systems* (pp. 2203–2213).
- Li, Z., Wang, C., Zheng, H., Zhang, J., & Li, B. (2022). FakeCLR: Exploring contrastive learning for solving latent discontinuity in data-efficient GANs. In *European conference on computer vision* (pp. 598–615).
- Lim, J. H., & Ye, J. C. (2017). Geometric GAN. arXiv:1705.02894, .
- Liu, B., Zhu, Y., Song, K., & Elgammal, A. (2021). Towards faster and stabilized GAN training for high-fidelity few-shot image synthesis. In *International conference on learning representations*.
- Mescheder, L. M., Geiger, A., & Nowozin, S. (2018). Which training methods for GANs do actually converge? In *International conference on machine learning* (pp. 3478–3487).
- Misra, I., & van der, M. L. (2020). Self-supervised learning of pretext-invariant representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 6707–6717).
- Miyato, T., Kataoka, T., Koyama, M., & Yoshida, Y. (2018). Spectral normalization for generative adversarial networks. In *International conference on learning representations*.
- Mo, S., Cho, M., & Shin, J. (2020). Freeze the discriminator: A simple baseline for fine-tuning GANs. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshop*.
- Nilsback, M., & Zisserman, A. (2008). Automated flower classification over a large number of classes. In *Proc. indian conf. comput. vis. graph. image process.* (pp. 722–729).
- Noguchi, A., & Harada, T. (2019). Image generation from small datasets via batch statistics adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 2750–2758).
- Nowozin, S., Cseke, B., & Tomioka, R. (2016). F-GAN: Training generative neural samplers using variational divergence minimization. In *Advances in neural information processing systems* (pp. 271–279).
- Oquab, M., Bottou, L., Laptev, I., & Sivic, J. (2014). Learning and transferring mid-level image representations using convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 1717–1724).
- Park, D., Kim, S., Lee, S., & Kim, H. J. (2024). Ddmi: Domain-agnostic latent diffusion models for synthesizing high-quality implicit neural representations. arXiv preprint arXiv:2401.12517, .
- Park, T., Liu, M., Wang, T., & Zhu, J. (2019). Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 2337–2346).
- Pope, P., Zhu, C., Abdelkader, A., Goldblum, M., & Goldstein, T. (2021). The intrinsic dimension of images and its impact on learning. In *International conference on learning representations*.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., & Sutskever, I. (2021). Learning transferable visual models from natural language supervision. In *International conference on machine learning* (pp. 8748–8763).
- Radford, A., Metz, L., & Chintala, S. (2016). Unsupervised representation learning with deep convolutional generative adversarial networks. In *International conference on learning representations*.
- Raghu, M., Unterthiner, T., Kornblith, S., Zhang, C., & Dosovitskiy, A. (2021). Do vision transformers see like convolutional neural networks? In *Advances in neural information processing systems* (pp. 12116–12128).
- Rangwani, H., Bansal, L., Sharma, K., Karmali, T., Jampani, V., & Babu, R. V. (2023). Noisytwins: Class-consistent and diverse image generation through stylegans. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 5987–5996).
- Richter, S. R., Alhaja, H. A., & Koltun, V. (2023). Enhancing photorealism enhancement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2), 1700–1715.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 10684–10695).
- Roth, K., Lucchi, A., Nowozin, S., & Hofmann, T. (2017). Stabilizing training of generative adversarial networks through regularization. In *Advances in neural information processing systems* (pp. 2018–2028).
- Saenko, K., Kulis, B., Fritz, M., & Darrell, T. (2010). Adapting visual category models to new domains. In *European conference on computer vision* (pp. 213–226).
- Salimans, T., Goodfellow, I. J., Zaremba, W., Cheung, V., Radford, A., & Chen, X. (2016). Improved techniques for training GANs. In *Advances in neural information processing systems* (pp. 2226–2234).
- Sauer, A., Chitta, K., Müller, J., & Geiger, A. (2021). Projected GANs converge faster. In *Advances in neural information processing systems* (pp. 17480–17492).
- Sauer, A., Schwarz, K., & Geiger, A. (2022). StyleGAN-XL: Scaling styleGAN to large diverse datasets. In *Proc. ACM SIGGRAPH conf* (pp. 1–10).
- Sharma, A., & Hamarneh, G. (2019). Missing MRI pulse sequence synthesis using multi-modal generative adversarial network. *IEEE Transactions on Medical Imaging*, 39(4), 1170–1183.
- Shin, Y., Byeon, Y., Son, G., Jang, H., Hwang, D., & Kim, S. (2025a). Physics-driven signal regularization in diffusion models for multi-contrast MR image synthesis. In *International conference on medical image computing and computer-assisted intervention* (pp. 403–413). Springer.
- Shin, Y., Lee, Y., Jang, H., Son, G., Kim, H., & Hwang, D. (2025b). Anatomical consistency and adaptive prior-informed transformation for multi-contrast MR image synthesis via diffusion model. In *Proceedings of the computer vision and pattern recognition conference* (pp. 30918–30927).
- Si, Z., & Zhu, S. (2012). Learning hybrid image templates (HIT) by information projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(7), 1354–1367.
- Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. In *International conference on learning representations*.
- Song, K., Xie, J., Zhang, S., & Luo, Z. (2023). Multi-mode online knowledge distillation for self-supervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 11848–11857).
- Sun, S., Luan, Z., Zhao, Z., Luo, S., & Han, S. (2024). Clr-gan: Improving gans stability and quality via consistent latent representation and reconstruction. In *European conference on computer vision* (pp. 210–227). Springer.
- Sungatullina, D., Zakharov, E., Ulyanov, D., & Lempitsky, V. S. (2018). Image manipulation with perceptual discriminators. In *European conference on computer vision* (pp. 587–602). (vol. 11210).
- Tan, M., & Le, Q. V. (2019). Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning* (pp. 6105–6114).
- Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., & Isola, P. (2020). What makes for good views for contrastive learning? In *Advances in neural information processing systems* (pp. 6827–6839).
- Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., & Jégou, H. (2021). Training data-efficient image transformers & distillation through attention. In *International conference on machine learning* (pp. 10347–10357).
- Tran, N., Tran, V., Nguyen, N., Nguyen, T., & Cheung, N. (2021). On data augmentation for GAN training. *IEEE Transactions on Image Processing*, 30, 1882–1897.
- Tran, N.-T., Tran, V.-H., Nguyen, B.-N., Yang, L., & Cheung, N.-M. (2019). Self-supervised gan: Analysis and improvement with multi-class minimax game. *Advances in Neural Information Processing Systems*, 32.
- Tseng, H., Jiang, L., Liu, C., Yang, M., & Yang, W. (2021). Regularizing generative adversarial networks under limited data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 7921–7931).
- Wang, T., Liu, M., Zhu, J., Tao, A., Kautz, J., & Catanzaro, B. (2018a). High-resolution image synthesis and semantic manipulation with conditional GANs. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 8798–8807).
- Wang, Y., Gonzalez-Garcia, A., Berga, D., Herranz, L., Khan, F. S., & van de Weijer, J. (2020). MineGAN: Effective knowledge transfer from GANs to target domains with few images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 9329–9338).
- Wang, Y., Mao, J., Guo, C., & Chen, S. (2024). Contrastive prototype-guided generation for generalized zero-shot learning. *Neural Networks*, 176, 106324.
- Wang, Y., Wu, C., Herranz, L., van de Weijer, J., Gonzalez-Garcia, A., & Raducanu, B. (2018b). Transferring GANs: Generating images from limited data. In *European conference on computer vision* (pp. 220–236).
- Wang, Z., Jiang, Y., Zheng, H., Wang, P., He, P., Wang, Z., Chen, W., Zhou, M. et al. (2023a). Patch diffusion: Faster and more data-efficient training of diffusion models. *Advances in Neural Information Processing Systems*, 36, 72137–72154.
- Wang, Z., Zheng, H., He, P., Chen, W., & Zhou, M. (2023b). Diffusion-GAN: Training GANs with diffusion. In *International conference on learning representations*.
- Yang, C., Shen, Y., Xu, Y., Zhao, D., Dai, B., & Zhou, B. (2022). Improving gans with a dynamic discriminator. *Advances in Neural Information Processing Systems*, 35, 15093–15104.
- Yang, C., Shen, Y., Xu, Y., & Zhou, B. (2021). Data-efficient instance generation from instance discrimination. In *Advances in neural information processing systems* (pp. 9378–9390).

- Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. (2014). How transferable are features in deep neural networks? In *Advances in neural information processing systems* (pp. 3320–3328).
- Yu, F., Seff, A., Zhang, Y., Song, S., Funkhouser, T., & Xiao, J. (2015). Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. arXiv:1506.03365, .
- Yu, N., Liu, G., Dundar, A., Tao, A., Catanzaro, B., Davis, L., & Fritz, M. (2021). Dual contrastive loss and attention for GANs. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 6711–6722).
- Zbontar, J., Jing, L., Misra, I., LeCun, Y., & Deny, S. (2021). Barlow twins: Self-supervised learning via redundancy reduction. In *International conference on machine learning* (pp. 12310–12320).
- Zeiler, M. D., & Fergus, R. (2014). Visualizing and understanding convolutional networks. In *European conference on computer vision* (pp. 818–833).
- Zhang, D., & Khoreva, A. (2019). Progressive augmentation of GANs. In *Advances in neural information processing systems* (pp. 6246–6256).
- Zhang, H., Chen, B., Tian, L., Wang, Z., & Zhou, M. (2020a). Variational hetero-encoder randomized GANs for joint image-text modeling. In *International conference on learning representations*.
- Zhang, H., Cissé, M., Dauphin, Y. N., & Lopez-Paz, D. (2018a). Mixup: Beyond empirical risk minimization. In *International conference on learning representations*.
- Zhang, H., Goodfellow, I. J., Metaxas, D. N., & Odena, A. (2019). Self-attention generative adversarial networks. In *International conference on machine learning* (pp. 7354–7363).
- Zhang, H., Xu, T., & Li, H. (2017). StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE/CVF international conference on computer vision* (pp. 5908–5916).
- Zhang, H., Zhang, Z., Odena, A., & Lee, H. (2020b). Consistency regularization for generative adversarial networks. In *International conference on learning representations*.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., & Wang, O. (2018b). The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 586–595).
- Zhang, Z., Hua, Y., Sun, G., Wang, H., & McLoone, S. (2024). Improving the training of the GANs with limited data via dual adaptive noise injection. In *Proceedings of the 32nd ACM international conference on multimedia* (pp. 6725–6734).
- Zhang, Z., Hua, Y., Sun, G., Wang, H., & McLoone, S. (2025). Improving the training of data-efficient GANs via quality aware dynamic discriminator rejection sampling. In *Proceedings of the computer vision and pattern recognition conference* (pp. 30682–30691).
- Zhao, M., Cong, Y., & Carin, L. (2020a). On leveraging pretrained GANs for generation with limited data. In *International conference on machine learning* (pp. 11340–11351).
- Zhao, S., Liu, Z., Lin, J., Zhu, J., & Han, S. (2020b). Differentiable augmentation for data-efficient GAN training. In *Advances in neural information processing systems* (pp. 7559–7570).
- Zhao, Z., Zhang, Z., Chen, T., Singh, S., & Zhang, H. (2020c). Image augmentations for GAN training. arXiv:2006.02595, .