

PetsGAN: Rethinking Priors for Single Image Generation

Zicheng Zhang¹, Yinglu Liu², Congying Han^{1*}, Hailin Shi², Tiande Guo¹, Bowen Zhou²

¹ University of Chinese Academy of Sciences

² JD AI Research

zhangzicheng19@mailsucas.ac.cn, liuyinglu1@jd.com, hancy@ucas.ac.cn

shihailin@jd.com, tdguo@ucas.ac.cn, bowen.zhou@jd.com

Abstract

Single image generation (SIG), described as generating diverse samples that have similar visual content with the given single image, is first introduced by SinGAN which builds a pyramid of GANs to progressively learn the internal patch distribution of the single image. It also shows great potentials in a wide range of image manipulation tasks. However, the paradigm of SinGAN has limitations in terms of generation quality and training time. Firstly, due to the lack of high-level information, SinGAN cannot handle the object images well as it does on the scene and texture images. Secondly, the separate progressive training scheme is time-consuming and easy to cause artifact accumulation. To tackle these problems, in this paper, we dig into the SIG problem and improve SinGAN by fully-utilization of internal and external priors. The main contributions of this paper include: 1) We introduce to SIG a regularized latent variable model. To the best of our knowledge, it is the first time to give a clear formulation and optimization goal of SIG, and all the existing methods for SIG can be regarded as special cases of this model. 2) We design a novel Prior-based end-to-end training GAN (PetsGAN) to overcome the problems of SinGAN. For one thing, we inject external priors obtained by the GAN inversion with PetsGAN to alleviate the problem of lack of high-level information for generating natural, reasonable and diverse samples, even for the object images. For another, we inject the internal priors of patch distribution correction with PetsGAN to ease and speed up the patch distribution learning. Moreover, our method gets rid of the time-consuming progressive training scheme and can be trained end-to-end. 3) We construct abundant qualitative and quantitative experiments to show the superiority of our method on both generated image quality, diversity, and the training speed. Moreover, we apply our method to other image manipulation tasks (*e.g.*, style transfer, harmonization), and the results further prove the effectiveness and efficiency of our method.

Introduction

The ability of modern deep learning to fulfill computer vision tasks is closely related to the building of large training

*Corresponding author

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This work was conducted when Zicheng Zhang was doing internship at JD AI Research.

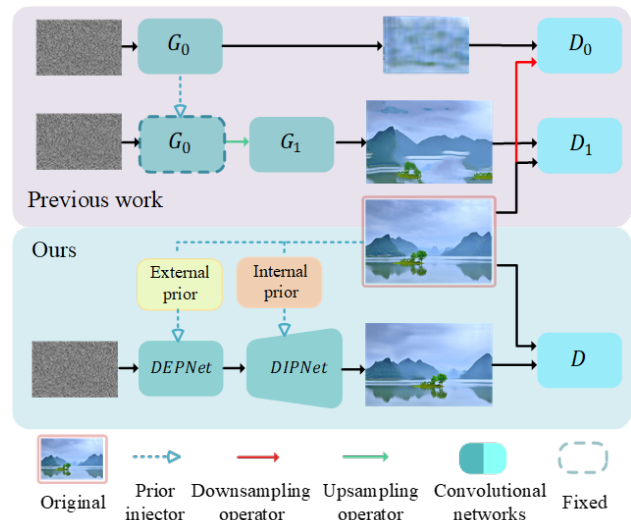


Figure 1: Comparing PetsGAN with previous work. The upper part shows the paradigm of two-stage SinGAN, where the generator G_0 is trained on low resolution images first, then copied to the next stage with fixed weights for high resolution generation. In contrast, PetsGAN is trained in one-stage by making full use of external and internal priors, which embodies significant advantages in both training efficiency and generated image quality and diversity. See more details in Sec. .

sets (Deng et al. 2009; Yu et al. 2015). It is, however, costly and difficult to collect training data in some scenarios, *e.g.*, medical diagnosis, industrial defects. Generative models, such as generative adversarial networks (GANs) (Goodfellow et al. 2014) and variational autoencoder (VAE) (Kingma and Welling 2013), have been expected to augment the training dataset, but all previous and current state-of-the-art models like DCGAN (Radford, Metz, and Chintala 2016), BigGAN (Brock, Donahue, and Simonyan 2019) and StyleGAN (Karras, Laine, and Aila 2019; Karras et al. 2020), face the same chicken-and-egg problem, which means that training generative models also requires a lot of data. As a new research direction that fully exploits the internal statistics of natural signals by deep learning, deep internal learning

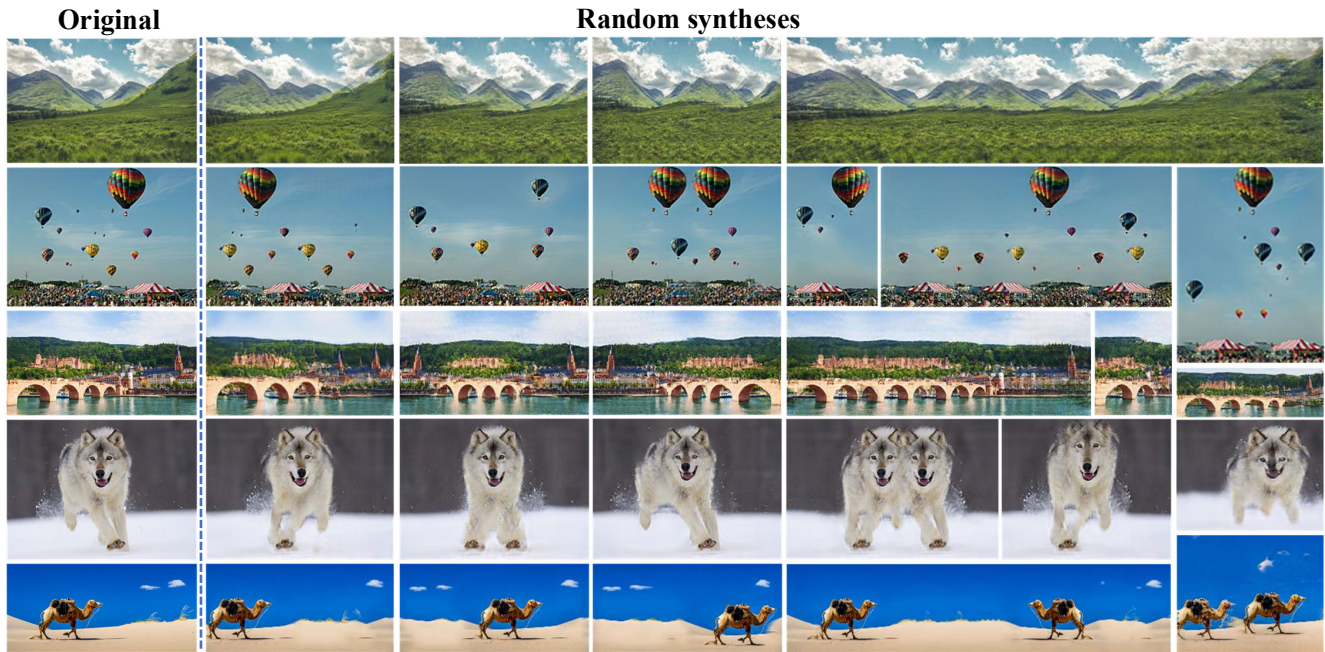


Figure 2: The first column refers to the original images, and the rest are random syntheses by PetsGAN. We can observe that PetsGAN can synthesize samples with arbitrary sizes, reasonable structure, fine texture, and great visual diversity.

(Ulyanov, Vedaldi, and Lempitsky 2017; Shocher, Cohen, and Irani 2018) has taken off and attracted wide attention since it only requires test image itself without additional supervision.

Not long ago deep internal learning has surprisingly stepped into the task of single image generation. Different from and beyond single image texture generative models (Jetchev, Bergmann, and Vollgraf 2016; Zhou et al. 2018) that synthesize images with repeated texture patterns, single image generation models (SIG models) (Shaham, Dekel, and Michaeli 2019; Shocher et al. 2019) are learned from the natural image and generate *similar* but *different* samples from the given image, where the former means similar visual content and the latter means diverse layout and appearance. The pioneering SinGAN (Shaham, Dekel, and Michaeli 2019) is the first unconditional SIG model, which builds a Laplacian pyramid of GANs with fully convolutional generators and patch discriminators to learn different scale information within the single image. The impressive performance of SinGAN proves the feasibility of internal learning on the generation task. From then on, many SIG models (Hinz et al. 2021; Chen et al. 2021; Zhang, Han, and Guo 2021; Sushko, Gall, and Khoreva 2021; Granot et al. 2021) are proposed to explore along this line.

Although great breakthroughs and achievements have been made in this field, there are still many problems to be solved. Firstly, due to only acquiring the internal texture information of the given image, previous SIG models usually generate chaotic structures, which is especially serious on object images because the semantics and structural relationship between object parts are very important to synthesize plausible images. To alleviate this problem, (Wang et al.

2021a) utilize a pre-trained classification model to provide semantic information for image generation. However, the supervision is too coarse to guide high-fidelity image generation, and the model generalization is limited because 1) each generated sample should be optimized separately and 2) different classification models are needed for different image types. Secondly, almost all SIG models (Shaham, Dekel, and Michaeli 2019; Hinz et al. 2021; Chen et al. 2021; Zhang, Han, and Guo 2021) follow the paradigm of SinGAN (Shaham, Dekel, and Michaeli 2019) to equip the pyramid of generators and patch discriminators for internal patch learning. Although the progressive or cascaded learning of pyramid networks (Denton et al. 2015; Karras et al. 2017) reduces the difficulties of model optimization, the training is time-consuming and easy to cause artifacts accumulation.

In this paper, we concentrate on the above-mentioned prominent problems in the SIG task, and make contributions as follows:

1) We dive into the SIG problem in essence, that is, to learn a diverse distribution from the Dirac distribution composed of a single image. Obviously, it is an ill-posed restoration problem, and appropriate priors and regularization terms are the keys to solve it. Therefore, we construct a unified regularized latent variable model to formulate the SIG task. To the best of our knowledge, it is the first time to give a clear mathematical definition and optimization goal of this task. All the existing SIG models can be regarded as special cases of this model.

2) We design a novel Prior-based end-to-end training single image GAN (PetsGAN), which is infused with internal priors and external priors to overcome the problems of SinGAN and the derived methods (Fig. 1). For one thing, we

inject external priors obtained by the GAN inversion with a DEPNet to alleviate the problem of lacking high-level information for generating natural, reasonable and diverse samples, even for the object images. For another, we inject the internal priors of patch distribution correction with a DIPNet to ease and speed up the patch distribution learning. Benefiting from the fully-utilization of external and internal priors, our method gets rid of the time-consuming progressive training scheme and can be trained end-to-end.

3) We construct abundant qualitative and quantitative experiments on a variety of datasets. Fig. 2 shows some syntheses of PetsGAN, experimental comparisons results show our method significantly surpasses other SIG models on image quality, diversity, and training speed. Moreover, we apply our method to other applications (*e.g.*, fast training of high resolution image generation, style transfer, harmonization). The results further prove the effectiveness and efficiency of our method.

Related Work

Single image generation model. Begin with modeling the internal patch distribution of a single image, SIG models aim to generate new plausible images with high texture quality. InGAN (Shocher et al. 2019) first proposes the conditional SIG model with the encoder-decoder architecture to make the geometric transformation of image look more natural. SinGAN (Shaham, Dekel, and Michaeli 2019) first proposes the unconditional SIG model composed of a Laplacian pyramid of GANs with fully convolutional generators and patch discriminators to learn different scale information. ConSinGAN (Hinz et al. 2021) modifies the training strategy of PGGAN (Karras et al. 2017) to improve SinGAN by training multiple stages simultaneously. PatchGenCN (Zheng, Xie, and Li 2021) proposes a pyramid of energy functions to explicit express the distribution within a single natural image. GPNN (Granot et al. 2021) replaces each stage of SinGAN with a non-training patch match module (Barnes et al. 2009) to speed up the generation. At the same time, a lot of work is no longer limited to simple texture generation. MOGAN (Chen et al. 2021) follows SinGAN but synthesizes the hand-marked regions of interest and the rest of the image separately, and then merges them into an unbroken image. One-shot GAN (Sushko, Gall, and Khoreva 2021) is an end-to-end model with multiple discriminators for learning different features of image, but it is not fully convolutional and more like normal GANs with diversity regularization (Yang et al. 2019). ExSinGAN (Zhang, Han, and Guo 2021) introduces GAN inversion (Pan et al. 2020) and perceptual loss (Johnson, Alahi, and Fei-Fei 2016) into SinGAN to improve performance on non-texture images.

Deep priors. In image restoration, priors mean some empirical knowledge or objective laws, which characterize the properties of the image and restrict the solution space. It can be explicitly expressed as a regularization term, or implicitly modeled in the loss function. For example, natural images are usually smooth in a local region (Buades, Coll, and Morel 2005; Mei et al. 2008), thus the smoothness can be taken as a prior in restoration tasks by total variation regular-




Figure 3: IR: image restoration network. PT: patch transfer. The scatters are the patches of low (left) and high (right) resolution, visualized by t-SNE. Red, blue and cyan represent the original, synthesis from DEPNet, and patch-transfer synthesis. The “star” represents the mean value. The distribution of cyan data can cover red data well, while a conspicuous deviation (caused by foreground patches) can be found between the blue data and red data.

ization (Rudin, Osher, and Fatemi 1992). Deep image prior (Ulyanov, Vedaldi, and Lempitsky 2018) demonstrates that the randomly initialized deep convolutional network can be miraculously considered as an implicit regularization term, which performs much better than the traditional regularization terms in restoration tasks. From then on, various deep priors (Gandelsman, Shocher, and Irani 2019; Gu, Shen, and Zhou 2020; Wang et al. 2021b) have been widely studied and applied in various fields. Deep generative prior (DGP) (Pan et al. 2020) assumes that the given image is sampled from the distribution space of the pre-trained generator, thus the generator is taken as an implicit regularization terms to restore image. IMAGINEA (Wang et al. 2021a) adopts the similar idea that the pre-trained classifier is utilized as the implicit regularization terms to exploit the semantics knowledge of a given image.

Methodology

Preliminaries

Notations. To simplify the expression, we assume that all data of d dimension comes from the set \mathcal{X}^d , and a measure space $(\mathcal{X}^d, \Sigma^d, \mu)$ can be defined over \mathcal{X}^d , where Σ^d is a Sigma algebra and μ is a measure. $\text{Prob}(\mathcal{X}^d)$ is the space of probability measures defined on \mathcal{X}^d . $\mathbb{P}_{real} \in \text{Prob}(\mathcal{X}^d)$ represents the real distribution of data, which is form-agnostic but there exists a sample set $\mathcal{D} = \{\mathbf{x}_i \in \mathcal{X}^d\}_{i=1}^n$. $\delta_{\mathbf{x}}$ represents the Dirac measure/distribution concentrated at \mathbf{x} . ρ denotes a generalized metric measuring the statistical distance between distributions. $\text{supp}(\mathbb{P}) := \{\mathbf{x} \in \mathcal{X}^d \mid \forall N_{\mathbf{x}} \in \Sigma^d, \mathbb{P}(N_{\mathbf{x}}) > 0\}$ denotes the support of distribution \mathbb{P} , where $N_{\mathbf{x}}$ is the neighborhood of \mathbf{x} .

Generalized Generation problem. Generation task aims at learning a parameterized distribution \mathbb{P}_{θ} to approximate the real distribution \mathbb{P}_{real} through the empirical distribution $\mathbb{P}_{\mathcal{D}}$,

$$\min_{\theta} \rho(\mathbb{P}_{\mathcal{D}}, \mathbb{P}_{\theta}). \quad (1)$$

There are extensive studies on problem (1), *e.g.*, the explicit density models (Kingma and Welling 2013; Rezende and Mohamed 2015; van den Oord, Kalchbrenner, and

Kavukcuoglu 2016) based on maximum log-likelihood estimation in which ρ is formed by KullbackLeibler (KL) divergence, and implicit density models like GANs in which ρ are formed by Jensen-Shannon (JS) divergence (Goodfellow et al. 2014), Wasserstein distance (Arjovsky, Chintala, and Bottou 2017), etc.

Patch Distribution. Suppose that $1 \leq j, s \leq d$ represent index and size of sliding window, respectively. For $\forall \mathbf{x} = [x_1, \dots, x_d]^T \in \mathcal{X}$, its j -th patch $\omega_j^{x,s} = [x_{\text{mod}(j-\lfloor s/2 \rfloor, d)}, \dots, x_{\text{mod}(j+\lfloor s/2 \rfloor, d)}]^T$. Denoting the mapping $\pi : \Sigma^d \rightarrow \Sigma^s, \mathcal{D} \rightarrow \mathcal{D}^s := \bigcup_{\mathbf{x} \in \mathcal{D}} \{\omega_j^{x,s}\}_{j=1}^d$. The random variable ω^s defined on dataset \mathcal{D}^s also has the empirical distribution $\mathbb{P}_{\mathcal{D}^s}$, which is called the *patch distribution* of \mathcal{D} with size s . Furthermore, π induces the mapping $\tau : \text{Prob}(\mathcal{X}^d) \rightarrow \text{Prob}(\mathcal{X}^s), \mathbb{P}_{\mathcal{D}} \rightarrow \mathbb{P}_{\mathcal{D}^s}$, and $\rho_\tau := \rho(\tau(\cdot), \tau(\cdot))$ is the induced metric for patch distribution.

Single Image Generation Problem

The SIG problem is first roughly described by (Shaham, Dekel, and Michaeli 2019) as generating diverse samples that have the same visual content as the given natural image. However, there is no clear mathematical definition so far. To better interpret and solve this problem, we introduce a regularized latent variable model and all the existing methods for SIG can be regarded as special cases of this model.

Problem statement. Given an image \mathbf{I} , the SIG problem is to learn a non-degenerate and realistic distribution \mathbb{P}_θ to restore the real distribution from the empirical distribution $\mathbb{P}_{\mathcal{D}}$. It’s worth noting that $\mathbb{P}_{\mathcal{D}}$ degenerates into a single-point distribution $\delta_{\mathbf{I}}$ (i.e., Dirac distribution) when $\mathcal{D} = \{\mathbf{I} \in \mathcal{X}^d\}$. Here “Non-degenerate” means $|\text{supp}(\mathbb{P}_{\text{real}})| > 1$ and “realistic” satisfies $\text{supp}(\mathbb{P}_\theta) \subseteq \text{supp}(\mathbb{P}_{\text{real}})$. Obviously, SIG is an ill-posed problem. In order to optimize it, appropriate regularization terms are needed to obtain the non-trivial solutions. Thus, the objective function can be defined by:

$$\min_{\theta} \rho(\delta_{\mathbf{I}}, \mathbb{P}_\theta) + \varphi(\mathbb{P}_\theta), \quad (2)$$

where ρ denotes the metric that measures the two distributions and φ denotes the regularization term that contains prior knowledge to make \mathbb{P}_θ diverse and realistic.

Regularized latent variable model. However, objective (2) is intractable due to the high dimensionality of \mathbf{I} and the difficulty of correctly choosing φ . A more efficient way is to introduce a lower dimensional latent random variable $\mathbf{c} \in \mathcal{C}$ to control \mathbf{x} . Thus, the objective can be changed to learn the joint distribution $\mathbb{P}_\theta(\mathbf{c}, \mathbf{x}) = \mathbb{P}_{\theta_0}(\mathbf{c})\mathbb{P}_{\theta_1}(\mathbf{x}|\mathbf{c})$, which is easier to solve by considering $\mathbb{P}(\mathbf{c})$ and $\mathbb{P}(\mathbf{x}|\mathbf{c})$ respectively. Given the sample \mathbf{I} and the corresponding latent code $\mathbf{c}_{\mathbf{I}}$, the optimization objective is:

$$\min_{\theta} \rho(\delta_{\mathbf{c}_{\mathbf{I}}, \mathbf{I}}, \mathbb{P}_\theta) + \varphi(\mathbb{P}_\theta), \quad (3)$$

where $\delta_{\mathbf{c}_{\mathbf{I}}, \mathbf{I}}$ is also a Dirac distribution of $(\mathbf{c}_{\mathbf{I}}, \mathbf{I})$. Actually, all the previous SIG models including (Shocher et al. 2019; Shaham, Dekel, and Michaeli 2019; Gur, Benaim, and Wolf 2020; Hinz et al. 2021; Zheng, Xie, and Li 2021; Zhang, Han, and Guo 2021) can be applied by our formulation (3).

Here, we take the canonical two-stage SinGAN (Shaham, Dekel, and Michaeli 2019) as an example to illustrate the definition and disadvantages of previous works. As Fig. 1 shows, the first stage in SinGAN is to learn $\mathbb{P}_{G_0}(\mathbf{c})$, where \mathbf{c} is downsampled image of \mathbf{x} . The learned distribution \mathbb{P}_{G_0} determines the model capability in diversity. The second stage is to learn $\mathbb{P}_{G_1}(\mathbf{x}|\mathbf{c})$, which plays a role as super-resolution or restoration based on the generated latent code by G_0 . The whole process can be formulated as:

$$\begin{aligned} & \min_{G_0, G_1} \rho_\tau(\delta_{\mathbf{c}_{\mathbf{I}}}, \mathbb{P}_{G_0}) + \rho_\tau(\delta_{\mathbf{I}}, \mathbb{P}_{G_1}) + \varphi(\mathbb{P}_{G_0}) \\ \text{s.t. } \varphi(\mathbb{P}_{G_0}) &= \begin{cases} 0 & \text{if } G_0 = \arg \min_{G_0} \rho_\tau(\delta_{\mathbf{c}_{\mathbf{I}}}, \mathbb{P}_{G_0}) \\ \infty & \text{otherwise} \end{cases} \end{aligned} \quad (4)$$

The objective (4) reveals the success of SinGAN: Firstly, the metric ρ_τ defined on *patch distribution* is benefit for diverse layout and high-fidelity texture synthesis. Secondly, the regularization φ suggests to train low-resolution generator G_0 and high-resolution generator G_1 progressively, which eases the instability of training. However, the drawbacks of SinGAN are also exposed: 1) the first stage only focuses on learning the low-level information i.e., internal patches, lacking of the high-level semantic or structural information which is important to synthesize natural and plausible samples. 2) the cascaded learning not only takes a long time but also suffers from artifacts accumulation. These drawbacks make SinGAN perform poorly on the non-textured images, so do the derived models adopting the paradigm of SinGAN.

PetsGAN

To comprehensively tackle the SIG problem, we propose a Prior-based end-to-end training single image GAN (PetsGAN), which is composed of a Deep External Prior Network (DEPNet, denoted as G) and a Deep Internal Prior Network (DIPNet, denoted as F). The overall objective for PetsGAN is defined by:

$$\min_{G, F} \rho_\tau(\delta_{\mathbf{I}}, \mathbb{P}_{F \circ G}) + \varphi(\mathbb{P}_G) + \phi(\mathbb{P}_F), \quad (5)$$

where ρ_τ is the metric of patch distribution. φ and ϕ are regularization terms corresponding to the external priors and the internal priors, respectively. As we have mentioned above, the patch distribution learning is a crucial factor for the success of SinGAN. Therefore, we adopt the same strategy to minimize the patch adversarial loss between the generated images and the original image. However, different from SinGAN which employs multiple discriminators for different stages, we only utilize one multi-scale discriminator (denoted as D) with different dilated kernel sizes without gradient penalty, so that the training time is reduced significantly. Moreover, we design novel regularization terms to achieve high-diversity in layout and high-fidelity in texture by fully utilizing the external priors and internal priors. Next, we will introduce the two parts in details.

External priors for high-diversity and reasonable layout.

The lack of high-level semantic and structural information makes SinGAN difficult to synthesize diverse and plausible samples, especially for non-texture images (e.g., object

images). To alleviate this problem, we employ GAN inversion to inject external priors. GAN inversion has been well studied (Zhu et al. 2016; Pan et al. 2020; Huh et al. 2020) in recent years. Resorting to the pre-trained generator G_{pre} , image \mathbf{I} can be projected into the latent space and obtained an optimal latent code z^* , and then extensive samples can be synthesized by disturbing z^* with Δz . Although the disturbance may reduce the quality of texture, it can generate abundant syntheses covering a large variation in layout. We downsample these syntheses to compress the texture details and utilize them to train a lightweight Deep External Prior Network (DEPNet). In practice, we define $\varphi(\mathbb{P}_G)$ by

$$\begin{aligned} \varphi(\mathbb{P}_G) = & \mathbb{E}_{x \sim \mathbb{P}_{G_{pre}}, z \sim \mathcal{N}(\mathbf{0}, \mathbf{1})} [\log D_G(x) - \log D_G(G(z))] \\ & - \mathbb{E}_{z_1, z_2 \sim \mathcal{N}(\mathbf{0}, \mathbf{1})} |D_G(G(z_1)) - D_G(G(z_2))|, \end{aligned} \quad (6)$$

where G refers to DEPNet and D_G is discriminator. The first term in Eq.(6) is the non-saturating GAN loss (Goodfellow et al. 2014). The second term is the collapse penalty to prevent G from mode collapse. The noise is concatenated with the sinusoidal positional encoding as the input to enhance the positional bias of DEPNet (Xu et al. 2020).

Internal priors for fast restoration of high-fidelity texture. The remarkable trait of SIG models is that, it makes full use of patch discriminator (Li and Wand 2016; Isola et al. 2017) to keep the synthesized texture consistent with that of the original image, *i.e.*, minimizing $\rho_\tau(\delta_{\mathbf{I}}, \mathbb{P}_F) = \rho(\mathbb{P}_{\mathbf{I}}^s, \mathbb{P}_F^s)$. However, training patch discriminator is very tricky and time-consuming. Additional gradient regularization like gradient penalty (Gulrajani et al. 2017) or R1 regularization (Mescheder, Geiger, and Nowozin 2018) is required to moderate the gradient from discriminator. Since the patch distribution $\mathbb{P}_{\mathbf{I}}^s$ is fixed, are there some efficient ways to approach it? Here, we show that the trait can be theoretically realized effectively via the reconstruction loss.

Proposition 1. *Considering image \mathbf{I} and its downsampled version $\mathbf{c}_{\mathbf{I}}$, an image restoration network (IR) can be well designed such that $\mathbf{I} = IR(\mathbf{c}_{\mathbf{I}})$ and*

- (i) *For appropriate s , there exist $s^* \leq s$ such that IR determines a mapping from $\mathbb{P}_{\mathbf{c}_{\mathbf{I}}}^{s^*}$ to $\mathbb{P}_{\mathbf{I}}^s$.*
- (ii) *If $\rho(\mathbb{P}_{\mathbf{c}_{\mathbf{I}}}^{s^*}, \mathbb{P}_{\mathbf{c}}^{s^*}) = 0$, then $\rho(\mathbb{P}_{\mathbf{I}}^s, \mathbb{P}_{IR(\mathbf{c})}^s) = 0$.*

Proof is provided in Appendix. Proposition 1 (i) exhibits that IR trained by reconstruction loss not only realizes pixel-level recovery, but also acts as a bridge between the patch distribution of low resolution images and that of high resolution images. (ii) exposes an imaginative way to capture the trait of previous SIG models, *i.e.*, reducing the differences of patch distribution between the generated samples and the original image in low resolution.

We inject internal priors of patch distribution correction into DIPNet F to realize fast restoration from low resolution syntheses to high resolution syntheses. DIPNet is a fully convolutional network containing two sub-modules, respectively patch transfer module (PT) for replacing each patch of synthesis c with the nearest patch of original image, and IR module to restore the new synthesis. For the $\omega_i^{c_1, s}$, PT learns the corresponding patch $a_{i, i^*} \omega_{i^*}^{c_1, s}$, where a_{i, i^*} (Eq.

(7)) is the coefficient from the attention map $\mathbf{A} = (a_{i, j})_{d \times d}$ to make the patch match differentiable.

$$\begin{aligned} b_{i, j} &= \text{sim}(\text{Conv}(\omega_i^{c_1, s}), \text{Conv}(\omega_j^{c_1, s})) \\ (a_{i, j})_{j=1, \dots, d} &= \text{softmax}(b_{i, 1}, \dots, b_{i, d}), \\ i^* &= \arg \max_j a_{i, j}. \end{aligned} \quad (7)$$

The sim denotes the cosine similarity, and Conv is initialized by the pre-trained VGG (Simonyan and Zisserman 2015). The perceptual distance (Johnson, Alahi, and Fei-Fei 2016) computed in feature space is more robust than in the pixel space (Zhang et al. 2019). These patches are integrated as a new synthesis, where overlap regions between different patches are averaged, then the new synthesis is fed into IR to restore high resolution. In practice, we form $\phi(\mathbb{P}_F)$ as

$$\phi(\mathbb{P}_F) = \|F(\mathbf{c}_{\mathbf{I}} + \Delta \mathbf{c}) - \mathbf{I}\| - \|\mathbf{A}\|_F^2, \quad (8)$$

where $\Delta \mathbf{c}$ is a small disturbance to make DIPNet more robust, and a sparsity constraint is imposed on \mathbf{A} to make a_{i, i^*} close to 1 (Remark 1 in Appendix). Fig. 3 shows the strength of internal priors. Although IR can restore \mathbf{I} from $\mathbf{c}_{\mathbf{I}}$ well (red images), for $c \sim \mathbb{P}_G$, $IR(c)$ produces very poor texture (blue images). After patch transfer, $IR(PT(c))$ has similar texture to \mathbf{I} (cyan images). The phenomenon in Fig. 3 is in line with Proposition 1. In Sec. we also show its amazing ability to save training time on high resolution images.

In summary, the PetsGAN minimize the patch generative adversarial loss ρ_τ and two regularization terms φ and ϕ corresponding to external and internal priors. The joint training of DEPNet and DIPNet is conducive to mutual adaptation, and the ablation study also proves the effectiveness.

Experiments

In this section, we evaluate our method on four datasets and compare the results with other state-of-the-art SIG models, *i.e.*, DGP (Pan et al. 2020), SinGAN (Shaham, Dekel, and Michaeli 2019), ConSinGAN (Hinz et al. 2021), HPVAE-GAN (Gur, Benaim, and Wolf 2020), PatchGenCN (Zheng, Xie, and Li 2021) ExSinGAN (Zhang, Han, and Guo 2021), qualitatively and quantitatively.

Training Settings

Datasets. For comprehensively evaluating the generative capability of PetsGAN, various experimental images are collected from various sources. The first one is the Places50 dataset (Shaham, Dekel, and Michaeli 2019), which contains 50 natural landscape scene images, the second one is the LSUN50 dataset (Hinz et al. 2021), which contains 50 images with more complex scene images, and the third one is the ImageNet50 dataset (Zhang, Han, and Guo 2021), which contains 50 images of objects. Besides, we also qualitatively show syntheses of images from the internet to illustrate the generalization of PetsGAN.

Training Details. Our experiments are conducted on NVIDIA RTX 3090. The input image \mathbf{I} is resized to make the longer side no more than 256, then downsampled with scale factor 8 to obtain the low-resolution image $\mathbf{c}_{\mathbf{I}}$. We adopt DGP (Pan et al. 2020) as the inversion method, in

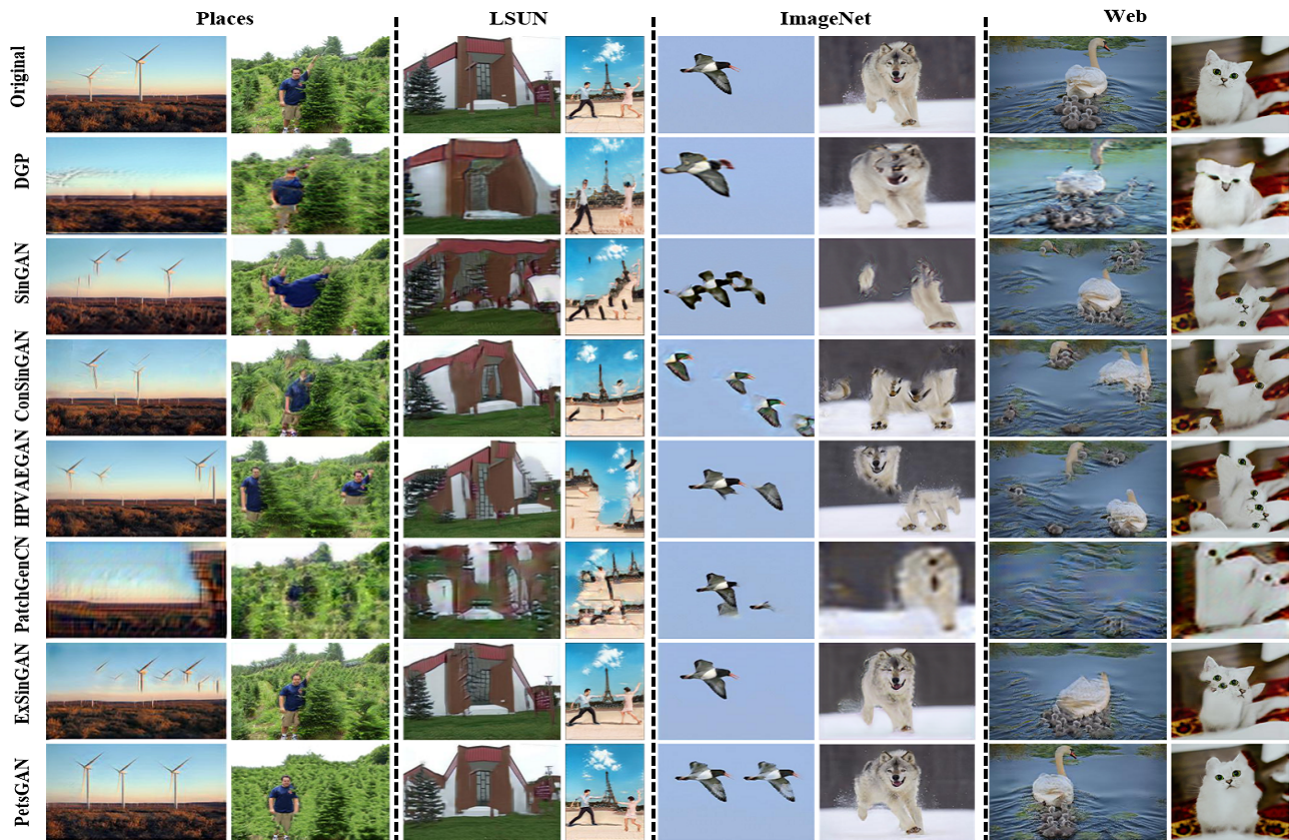


Figure 4: Visualization results of different SIG models on different datasets. Our method has great advantage in the quality of photo-realistic generation while guaranteeing the diversity, for not only scene but also object images.

which disturbance is sampled from $\mathcal{N}(0, 0.5)$. The input noise of PetsGAN is sampled from $\mathcal{N}(0, 1)$. The window size s is set to 7. We use Adam optimizer for G, F, D_G, D but with different learning rates. The training epoch of PetsGAN is 5000. The batch sizes for optimizing $\rho_+(\delta_{\mathbf{I}}, \mathbb{P}_{F \circ G})$ and $\varphi(\mathbb{P}_G)$ are 1 and 32, respectively. The warm up strategy is used to train G and F . Due to space limitation, the detailed network structures and learning strategy are provided in the Appendix.

Experimental Results

Qualitative analysis. Fig. 2 gives syntheses of PetsGAN on various exemplars with arbitrary sizes. We can see PetsGAN performs excellent on both texture and object images. **It can generate images with diverse layouts while the content and texture are consistent with the exemplars.** To demonstrate the superiority of PetsGAN, we compare PetsGAN with other SIG models. As Fig. 4 shows, SinGAN and ConSinGAN usually generate unnatural layouts or structures due to lack of semantic information, *e.g.*, windmills and human bodies. DGP utilizes the external information to generate reasonable structures but the texture is synthesized coarsely. ExSinGAN also embodies the issue of diversity because the semantic consistence constraint between the synthesized samples and the source image. By contrast, Pets-

GAN surpasses them in terms of both diversity and quality.

Quantitative analysis. We compare the performance of SIG models in terms of generative quality, diversity, and training time quantitatively. Similar to ExSinGAN, we adopt SIFID and LPIPS to measure the image quality and diversity, respectively. SIFID quantifies the distance between the generated image and the original image in feature space. LPIPS (Zhang et al. 2018) compares the perceptual differences using pre-trained AlexNet, which is insensitive to color changes but more sensitive to spatial changes. The measured values are provided in Tab. 1. We find that the similar conclusion can be drawn as the qualitative analysis. Although some methods achieve high scores in LPIPS, they perform poorly in SIFID, and the visualization results are bad in this case. Compared with other methods, PetsGAN achieves low scores in SIFID and high scores in LPIPS simultaneously, which means PetsGAN can generate diverse samples with natural layout and realistic texture. Besides, our model has significant advantages in training efficiency. As Fig. 5 shows, the average training time of our model is 10 minutes, almost 4 times faster than SinGAN.

Ablation study. To better evaluate the effect of training strategies and priors, we conduct a series of ablation studies (Table 1). Firstly, comparing the joint training with cas-

	Places50		LSUN50		ImageNet50	
	SIFID↓	LPIPS↑	SIFID↓	LPIPS↑	SIFID↓	LPIPS↑
DGP	0.64	0.28	0.75	0.36	1.65	0.31
SinGAN	0.09	0.25	0.23	0.33	0.60	0.33
ConSinGAN	0.06	0.24	0.11	0.32	0.56	0.39
HPVAEGAN	0.30	0.38	0.35	0.44	0.55	0.47
PatchGenCN	1.23	0.60	1.24	0.62	1.29	0.61
ExSinGAN	0.10	0.23	0.11	0.25	0.45	0.24
PetsGAN	0.08	0.31	0.09	0.33	0.21	0.35
cascaded	0.10	0.31	0.11	0.34	0.50	0.37
w/o $\rho_\tau(\delta_I, \mathbb{P}_{F \circ G})$	0.23	0.33	0.23	0.38	0.30	0.39
w/o $\varphi(\mathbb{P}_G)$	0.11	0.18	0.11	0.21	0.25	0.22
w/o $\phi(\mathbb{P}_F)$	0.20	0.23	0.19	0.26	0.40	0.33

Table 1: Quantitative results of different SIG models and ablation study. “cascaded” means training DEPNet individually, and then keeping it fixed and training DIPNet.




Figure 5: Training time of different SIG models.

caded training, the former’s SIFID is consistently smaller than the latter, conditioned on similar LPIPS, proving the joint training indeed helps DIPNet and DEPNet towards better mutual adaption. Secondly, patch adversarial loss is taken out, namely training DEPNet and DIPNet independently, and thereby the SIFID increases largely, proving its positive contributions to synthesis. Thirdly, if we do not use external priors, with plummeting LPIPS the syntheses lose the diversity largely, meaning that external prior is the key to provide the diversity, and avoid the overfitting. Finally, the SIFID and LPIPS both become worse if we do not use internal prior. This implies that the internal prior is not only helpful to the quality, but also prevents over-fitting to some extent. In summary, the ablation study demonstrated the rationality of PetsGAN’s design and the necessity of each part.

More Applications

Fast training of high resolution images. PetsGAN can also handle high resolution image generation by simple plugging in a new DIPNet, while keeping high efficiency in training. SinGAN takes about 6 hours for training with an image of size 1024×1024 on NVIDIA RTX 3090. The time-consuming training undoubtedly hinders the application prospect of SIG models. By contrast, we first obtain the 256×256 synthesis, and feed it into the new DIPNet with 4 times upscaling only trained by reconstruction loss. Thus, the overall training can be completed in less than 15 minutes, and the synthesized texture is of high-fidelity without



Figure 6: High resolution image generation.

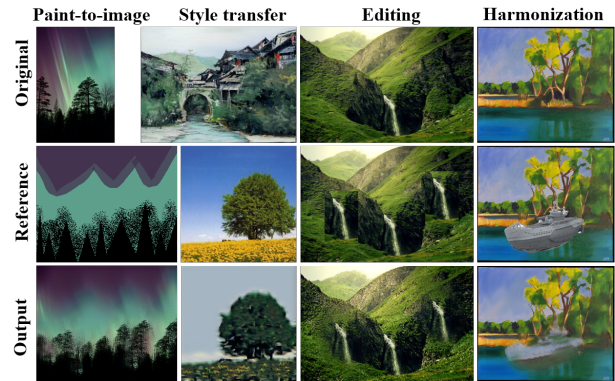


Figure 7: Applications to other image manipulation tasks.

degradation (Fig. 6) benefiting from the internal prior.

Image manipulation tasks. Our model can also be applied to many other image manipulation tasks, *e.g.*, paint-to-image, style transfer, editing and harmonization. A recent work SinIR (Yoo and Chen 2021) shows that the pyramid networks only with reconstruction loss can do these tasks effectively. Different from SinIR, we only use the one stage DIPNet with reconstruction loss to realize these applications. We show some results of our method in Fig. 7. The training time cost is less than 2 minutes, which is much faster than SinIR and other SIG models, and the generated results are also satisfactory.

Conclusion

In this paper, we focus on solving the SIG problem by fully-utilizing the prior knowledge. First, we formulate SIG with a regularized latent model, and analyze the shortcomings of SinGAN and other work from the form. Based on latent model, we propose PetsGAN using external and internal priors. The external priors are used to provide high-level information and prevent model overfitting. The internal prior is to reduce the patch distribution discrepancy between the syntheses and the input image, which can effectively speed up the training of the model. PetsGAN can achieve competitive results and improvements for SIG with more image types and faster training speed. It can also be applied to image manipulate tasks at a high speed.

Acknowledgements

This paper is supported by the National key research and development program of China (2021YFA1000403), the National Natural Science Foundation of China (Nos. U19B2040), the Strategic Priority Research Program of Chinese Academy of Sciences, Grant No. XDA27000000, and the Fundamental Research Funds for the Central Universities.

References

- Arjovsky, M.; Chintala, S.; and Bottou, L. 2017. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 214–223.
- Barnes, C.; Shechtman, E.; Finkelstein, A.; and Goldman, D. B. 2009. PatchMatch: a randomized correspondence algorithm for structural image editing. In *SIGGRAPH 2009*.
- Brock, A.; Donahue, J.; and Simonyan, K. 2019. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *International Conference on Learning Representations*.
- Buades, A.; Coll, B.; and Morel, J. 2005. A non-local algorithm for image denoising. *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, 2: 60–65 vol. 2.
- Chen, J.; Xu, Q.; Kang, Q.; and Zhou, M. 2021. MOGAN: Morphologic-structure-aware Generative Learning from a Single Image. *ArXiv*, abs/2103.02997.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Denton, E. L.; Chintala, S.; Szlam, A. D.; and Fergus, R. 2015. Deep Generative Image Models using a Laplacian Pyramid of Adversarial Networks. In *NIPS*.
- Gandelsman, Y.; Shocher, A.; and Irani, M. 2019. “Double-DIP”: Unsupervised Image Decomposition via Coupled Deep-Image-Priors. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11018–11027.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- Granot, N.; Shocher, A.; Feinsein, B.; Bagon, S.; and Irani, M. 2021. Drop the GAN: In Defense of Patches Nearest Neighbors as Single Image Generative Models. *ArXiv*, abs/2103.15545.
- Gu, J.; Shen, Y.; and Zhou, B. 2020. Image processing using multi-code gan prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3012–3021.
- Gulrajani, I.; Ahmed, F.; Arjovsky, M.; Dumoulin, V.; and Courville, A. C. 2017. Improved training of wasserstein gans. In *Advances in neural information processing systems*, 5767–5777.
- Gur, S.; Benaim, S.; and Wolf, L. 2020. Hierarchical Patch VAE-GAN: Generating Diverse Videos from a Single Sample. *ArXiv*, abs/2006.12226.
- Hinz, T.; Fisher, M.; Wang, O.; and Wermter, S. 2021. Improved techniques for training single-image gans. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1300–1309.
- Huh, M.; Zhang, R.; Zhu, J.-Y.; Paris, S.; and Hertzmann, A. 2020. Transforming and Projecting Images into Class-conditional Generative Networks. *arXiv preprint arXiv:2005.01703*.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-Image Translation with Conditional Adversarial Networks. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5967–5976.
- Jetchev, N.; Bergmann, U.; and Vollgraf, R. 2016. Texture synthesis with spatial generative adversarial networks. *arXiv preprint arXiv:1611.08207*.
- Johnson, J.; Alahi, A.; and Fei-Fei, L. 2016. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, 694–711. Springer.
- Karras, T.; Aila, T.; Laine, S.; and Lehtinen, J. 2017. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4401–4410.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and Improving the Image Quality of StyleGAN. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 8107–8116.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Li, C.; and Wand, M. 2016. Precomputed Real-Time Texture Synthesis with Markovian Generative Adversarial Networks. *ArXiv*, abs/1604.04382.
- Mei, T.; Yang, B.; Yang, S.; and Hua, X. 2008. Video collage: presenting a video sequence using a single image. *The Visual Computer*, 25: 39–51.
- Mescheder, L. M.; Geiger, A.; and Nowozin, S. 2018. Which Training Methods for GANs do actually Converge? In *ICML*.
- Pan, X.; Zhan, X.; Dai, B.; Lin, D.; Loy, C. C.; and Luo, P. 2020. Exploiting Deep Generative Prior for Versatile Image Restoration and Manipulation. *arXiv preprint arXiv:2003.13659*.
- Radford, A.; Metz, L.; and Chintala, S. 2016. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *CoRR*, abs/1511.06434.
- Rezende, D. J.; and Mohamed, S. 2015. Variational Inference with Normalizing Flows. In *ICML*.
- Rudin, L.; Osher, S.; and Fatemi, E. 1992. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60: 259–268.

- Shaham, T. R.; Dekel, T.; and Michaeli, T. 2019. SinGAN: Learning a Generative Model From a Single Natural Image. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 4569–4579.
- Shocher, A.; Bagon, S.; Isola, P.; and Irani, M. 2019. InGAN: Capturing and Retargeting the “DNA” of a Natural Image. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 4491–4500.
- Shocher, A.; Cohen, N.; and Irani, M. 2018. “zero-shot” super-resolution using deep internal learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3118–3126.
- Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR*, abs/1409.1556.
- Sushko, V.; Gall, J.; and Khoreva, A. 2021. One-Shot GAN: Learning to Generate Samples from Single Images and Videos. *ArXiv*, abs/2103.13389.
- Ulyanov, D.; Vedaldi, A.; and Lempitsky, V. 2017. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6924–6932.
- Ulyanov, D.; Vedaldi, A.; and Lempitsky, V. 2018. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9446–9454.
- van den Oord, A.; Kalchbrenner, N.; and Kavukcuoglu, K. 2016. Pixel Recurrent Neural Networks. *ArXiv*, abs/1601.06759.
- Wang, P.; Li, Y.; Singh, K. K.; Lu, J.; and Vasconcelos, N. 2021a. IMAGINE: Image Synthesis by Image-Guided Model Inversion. *ArXiv*, abs/2104.05895.
- Wang, X.; Li, Y.; Zhang, H.; and Shan, Y. 2021b. Towards Real-World Blind Face Restoration with Generative Facial Prior. *ArXiv*, abs/2101.04061.
- Xu, R.; Wang, X.; Chen, K.; Zhou, B.; and Loy, C. C. 2020. Positional Encoding as Spatial Inductive Bias in GANs. *ArXiv*, abs/2012.05217.
- Yang, D.; Hong, S.; Jang, Y.; Zhao, T.; and Lee, H. 2019. Diversity-Sensitive Conditional Generative Adversarial Networks. *ArXiv*, abs/1901.09024.
- Yoo, J.; and Chen, Q. 2021. SinIR: Efficient General Image Manipulation with Single Image Reconstruction. In *ICML*.
- Yu, F.; Seff, A.; Zhang, Y.; Song, S.; Funkhouser, T.; and Xiao, J. 2015. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhang, Z.; Han, C.; and Guo, T. 2021. ExSinGAN: Learning an Explainable Generative Model from a Single Image. *ArXiv*, abs/2105.07350.
- Zhang, Z.; Wang, Z.; Lin, Z.; and Qi, H. 2019. Image super-resolution by neural texture transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7982–7991.
- Zheng, Z.; Xie, J.; and Li, P. 2021. Patchwise Generative ConvNet: Training Energy-Based Models From a Single Natural Image for Internal Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2961–2970.
- Zhou, Y.; Zhu, Z.; Bai, X.; Lischinski, D.; Cohen-Or, D.; and Huang, H. 2018. Non-stationary texture synthesis by adversarial expansion. *ACM Transactions on Graphics (TOG)*, 37(4): 1–13.
- Zhu, J.-Y.; Krähenbühl, P.; Shechtman, E.; and Efros, A. A. 2016. Generative visual manipulation on the natural image manifold. In *European conference on computer vision*, 597–613. Springer.