

ExSinGAN: Learning an Explainable Generative Model From a Single Image

Zicheng Zhang
zhangzicheng19@mails.ucas.ac.cn

University of Chinese Academy of Science, Beijing, China

Congying Han*
hancy@ucas.ac.cn

Tiande Guo
tdguo@ucas.ac.cn

Abstract

Generating images from a single sample has attracted extensive attention recently. In this paper, we formulate the problem as sampling from the conditional distribution of a single image, and propose a hierarchical framework that simplifies the learning of the intractable conditional distribution through the successive learning of the distributions about structure, semantics, and texture, making the generative model more comprehensible compared with previous works. On this basis, we design ExSinGAN composed of three modular GANs for learning an explainable generative model from a given image, where the modular GANs model the distributions about structure, semantics, and texture successively. ExSinGAN is learned not only from the internal patches of the given image as the previous works did, but also from the external prior obtained by the GAN inversion technique. Benefiting from the appropriate combination of internal and external information, ExSinGAN has a more powerful capability of generation and competitive generalization ability for image manipulation tasks.

1 Introduction

Recently, single image generative models get more and more attention. The groundbreaking work [6] by Shaham *et al.* proposed a pyramid of fully convolutional GANs [5] named SinGAN to learn the internal distribution of patches within the single image. SinGAN is an unconditional generative model that acquires noise inputs and produces diverse syntheses but only needs the one-shot training image. Due to the impressive performance and extensible architecture, SinGAN has been the paradigm of single image generation, which is improved [16, 13, 19, 25, 22] and applied to other fields like video synthesis [18], structure analogy [4] and inpainting [7]. Although SinGAN has proved the feasibility of generating diverse images from a natural single image, experiments show that it is still difficult to obtain plausible syntheses when the given image is complex, *e.g.*, image with large objects. This drawback also exists in previous related works [8, 17, 15, 19, 36], which almost follow the paradigm of SinGAN to build the pyramid networks for internal learning [3].

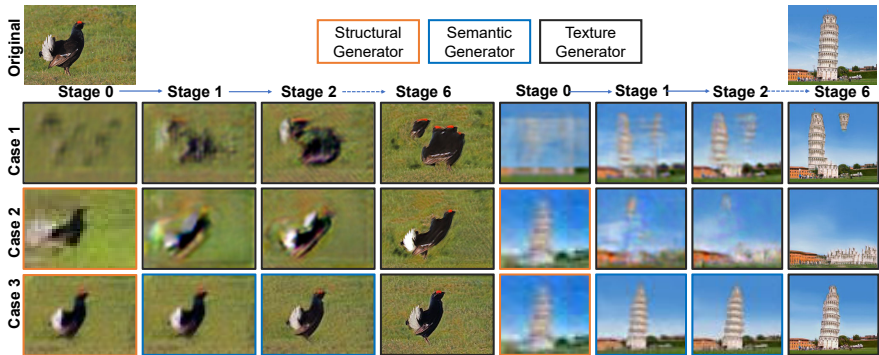


Figure 1: Each row illustrates the outputs from 7-stage single image generative models comprising generators in legend (Sec. 3 for details), and the border color of each output implies its source generator. All pictures have been upsampled to the same size as the original. Note that Case 1 and Case 3 stand for SinGAN [51] and our ExSinGAN respectively.

In this paper, we focus on improving the generative capability of SinGAN on arbitrary natural images, especially for those images including objects or complex scenes. Firstly, we revisit the core idea of SinGAN, namely the progressive learning strategy of pyramid networks. The progressive learning strategy [11, 21] training the pyramid networks stage by stage, has played an important role in stabilizing the performance of GANs. Specifically, for the large-scale progressively training model, the generator at the bottom learns to synthesize diverse but reasonable layouts of low resolution. In the subsequent stages, generators learn to add more semantic and texture details to the layouts. However, it is hard to obtain effective semantic and structural supervision from a single image. We observe that the bottom stage of SinGAN [51] generally synthesizes disorder layouts (Case 1 in Fig. 1), and the details replenished in the later stages are likely to be uncontrollable and meaningless. This phenomenon illustrates the difference between SinGAN and the large-scale progressive training model in nature, where each stage of SinGAN plays an unclear or even negative role in the generation. The entire model is more like a magic black box we can never predict how it will run and behave.

Our insight is that, both the reliability and generative capability of SinGAN can be improved by introducing additional structural and semantic supervision. In this paper, we propose a hierarchical framework for the single image generation problem and construct an *explainable single image generative model* called ExSinGAN. Specifically, ExSinGAN is learned from both internal patches and external priors of the given image via three modular GANs. Firstly, the structural GAN is devoted to synthesizing coarse but reasonable layout with the supervision of structural prior knowledge from GAN inversion [27]. Then, the semantic GAN is dedicated to injecting semantic details into the coarse layout with the supervision of a pre-trained classifier [20]. Finally, the texture GAN replenishes fine texture details to the previous output. To our best knowledge, ExSinGAN is the first single image generative model considering both internal information [51, 52] and external generative priors [27]. Through experiments (Sec. 4) we demonstrate that ExSinGAN can be adapted to a variety of images including texture, objects and scenes. Moreover, ExSinGAN also has competitive generalization capabilities on some image manipulation tasks like editing, harmonization and paint-to-image.

2 Related Work

Internal Learning Generative Model. Deep internal learning methods, which are trained by the internal information of a single image, have taken off and caught popular attention since ZSSR [52] and DIP [58]. InGAN [34] is the first internal learning generative model, which requires the pre-defined geometric transformation inputs to generate diverse images. SinGAN [60] is the first unconditional generative model, which comprises a pyramid of fully convolutional GANs to learn the internal patch distribution of the given image. SinGAN adopts the cascaded training strategy [10]: the bottom generator synthesizes the layout, and the later generators learn residual details to enlarge the layout. When the current generator is being trained, all lower stages generator are kept fixed. Some works have modified SinGAN from different perspectives. ConSinGAN [19] adopts the training strategy of PG-GAN [20] to improve the generative quality of SinGAN, and reduce the stages of the model to realize faster training. HP-VAEGAN [18] introduces patch-based VAE [23] to make the generation more diverse. MOGAN, which [8] follows the multi-stage architecture of SinGAN, firstly synthesizes the interesting region and the rest region of the image respectively, then merges them into an integrated image. GPNN [16] borrows the idea from the classical method PatchMatch [0], replacing each stage GAN with a non-training patch match module to generate a new image quickly. One-Shot GAN [56] is a one-stage model containing multiple discriminators to learn different features of the given image.

Generative Priors and Model Inversion. Different from the internal learning method, model inversion utilizes the pre-trained deep learning models as the external generative priors to conduct image synthesis. The most used pre-trained models include the classifiers [39] and GANs [40]. In this paper, we focus on the most popular way, GAN inversion. GAN inversion aims at finding a code in the latent space of a pre-trained GAN that best reconstructs the given image. Once the optimal latent code is found, new images can be synthesized by manipulating the latent code. The methods of GAN inversion can be divided into four categories. The first one based on gradient optimization minimizes the loss function directly [5, 9, 24, 28, 30, 33, 47]. The second one trains an encoder coupling with the generator to find the latent code indirectly [13, 14]. The third one combining the above two methods together [2, 26, 47], initializes the latent code via a coupled encoder, then optimizes the latent code using gradient optimization method. The last one is joint optimization [8, 27], which not only optimizes the latent code but also fine-tunes the parameters of the generator simultaneously. The representative model DGP [27] of joint optimization fine-tunes the generator of BigGAN [6] pre-trained on ImageNet [100] in a progressive manner, giving rise to more precise and faithful reconstruction for natural images.

3 Methodology

Preliminaries. Considering a single image $x \in \mathbb{R}^{H \times W}$ has size $H \times W$, a single image generative model formulates the conditional distribution $p(y|x)$, which we can sample from to get the synthesis $y \in \mathbb{R}^{H \times W}$. For convenience, we use $x^i \in \mathbb{R}^{H_i \times W_i}$ to denote the downsampled version of x , where the size of x^i increases with the stage i varying from 0 to N and $x = x_N$. Instead of learning the distribution $p(y|x)$, the progressive training of SinGAN aims at learning

$$p(y^0, y^1, \dots, y^N | x^0, x^1, \dots, x^N) \approx p(y^0 | x^0) \prod_{i=1}^N p(y^i | x^i, y^{i-1}), \quad (1)$$

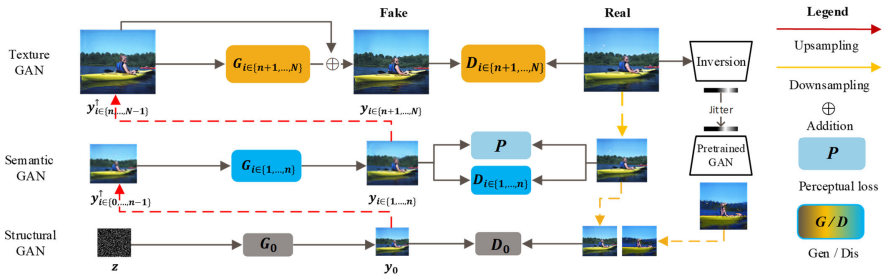


Figure 2: The architecture of ExSinGAN, which is composed of three modular GANs, the structural GAN for generating the coarse layout, semantic GAN for constraining the semantics of synthesis, the texture GAN for making the synthesis exquisite. Except for structural GAN, both semantic GAN and texture GAN are multi-stage. The GAN inversion provides a structural prior to train structural GAN.

where y^i can be sampled successively to get the final synthesis y^N . We make an assumption that image x can be decomposed into $\{x_{str}, x_{sem}, x_{tex}\}$ called the *triplet representation*, where x_{str} , x_{sem} and x_{tex} represent the structure, semantics and texture of x respectively.

We firstly point that such representation exists and $x_{str}, x_{sem}, x_{tex}$ can be well-defined on specific manifolds. Firstly, Considering a large-scale pre-trained generator G_{pre} and z^* subjected to $x = G_{pre}(z^*)$, we formulate $x_{str} := G_{pre}(z^*)^0$. The superscript ‘0’ represents the downsampling operator as defined in Eq. (1), which aims to compress the texture and keep the layout of x effectively. Note that even though x_{str} and x^0 are equal, x_{str} belongs to the new manifold defined by $G_{pre}(z), z \in$ latent space. Unlike that disturbance in pixel space usually brings noise to image, the disturbance to latent code z^* likely changes the structure of x reasonably, where the neighbors of x_{str} have similar structure to x_{str} . Then, the semantic feature x_{sem} can be extracted using pre-trained convolutional networks ϕ such as VGG19 [5], *i.e.*, $x_{sem} = \phi(x)$. Finally, the texture feature x_{tex} can be interpreted as the internal patches of x . Thus, we have determined the form of the triplet.

Proposed Method. It is not difficult to find that SinGAN just constrains the distance between y_{tex} and x_{tex} by patch discriminators, *i.e.*, maximizing $p(y_{tex}^0|x_{tex}^0) \prod_{i=1}^N p(y_{tex}^i|x_{tex}^i, y^{i-1})$, hence SinGAN hardly generates images with reasonable structure and semantics. According to triplet representation, we hope to maximize Eq. (2) to make all structure, semantics and texture reasonable. Therefore, we introduce the strengthening form of Eq. (1)

$$p(y_{str}|x_{str})p(y_{tex}^0|x_{tex}^0) \prod_{i=1}^n (p(y_{sem}^i|x_{sem}^i, y^{i-1})p(y_{tex}^i|x_{tex}^i, y^{i-1})) \prod_{i=n+1}^N p(y_{tex}^i|x_{tex}^i, y^{i-1}). \quad (2)$$

Eq. (2) consists of three parts, the first part $p(y_{str}|x_{str})p(y_{tex}^0|x_{tex}^0)$ makes the texture and structure of syntheses similar to the origin. The second part $\prod_{i=1}^n (p(y_{sem}^i|x_{sem}^i, y^{i-1})p(y_{tex}^i|x_{tex}^i, y^{i-1}))$ makes the texture and semantics of syntheses similar to the origin at the intermediate stages. and the third part $\prod_{i=n+1}^N p(y_{tex}^i|x_{tex}^i, y^{i-1})$ only makes the texture of syntheses similar to the origin in high resolution. Fig. 2 shows the overview of ExSinGAN, where the network comprises three modular GANs, respectively structural GAN, semantic GAN, and texture GAN aiming to model three parts in Eq.(2) successively. In the following section, we simply denote each generator and discriminator of ExSinGAN as G_i and D_i , $i \in \{0, \dots, N\}$.

Structural GAN. The structural GAN implicitly learns $p(y_{tex}^0|x_{tex}^0)p(y_{str}|x_{str})$ from the image x , where the generator G_0 synthesizes images of reasonable layouts and high qual-

ity of texture. We firstly adopt GAN inversion method DGP [27] to find z^* in the latent space of BigGAN [8], then disturb z^* via Gaussian noise to generate dataset $Data = \{G_{pre}(z^* + \Delta z_i)^0 | i = 1, \dots, m\}$, where Δz_i is a random noise and m is the number of syntheses. Moreover, because the texture and color of syntheses in $Data$ differ greatly from the original image, we insert some copies of x^0 into $Data$ to suppress these artifacts. D_0 is designed as a fully-connected discriminator to learn the texture and structure together, which is different from that SinGAN [5] uses a patch discriminator. WGAN-gp loss [17]

$$\min_{G_0} \max_{D_0} \mathcal{L}_{adv}(G_0, D_0) \quad (3)$$

is used as the adversarial loss to keep training stable. Sinusoidal positional encoding is concatenated with the input noise [12] to enhance the spatial inductive bias of G_0 [42]. Once the structural GAN is trained, G_0 can generate low-resolution layouts of higher quality compared with SinGAN (stage 0 in Fig. 1).

Semantic GAN. According to $\prod_{i=1}^n (p(y_{sem}^i | y_{sem}^{i-1}) p(y_{tex}^i | x_{tex}^i, y^{i-1}))$, semantic GAN aims at adding more texture and semantic details to the layout from G_0 . For the semantic details, we adopt the perceptual loss [24]

$$\mathcal{L}_p(x^i, G_i(y^{i-1\uparrow})) = \|\phi(x) - \phi(G_i(y^{i-1\uparrow}))\| \quad (4)$$

to enlarge x^0 meaningfully, where " \uparrow " means bicubic interpolation, and G_i is a fully convolutional network. Here we do not use residual learning like SinGAN, because it will cause serious artifacts in practice. Patch discriminator D_i is used for patch adversarial learning, and the reconstruction loss $\mathcal{L}_{rec}(G_i) = \|G_i(x^{i-1\uparrow}) - x^i\|^2$ aims to stabilize the training. Denoting $\mathcal{L}_p(x^i, G_i(y^{i-1\uparrow}))$ as $\mathcal{L}_p(G_i)$, the total training loss is

$$\min_{G_i} \max_{D_i} \mathcal{L}_{adv}(G_i, D_i) + \alpha_1 \mathcal{L}_{rec}(G_i) + \lambda \mathcal{L}_p(G_i). \quad (5)$$

For arbitrary y^0 from G_0 , the output y^n has similar layout to y^0 and close semantics to x , and without semantic generator the quality will be poor (Case 2 and Case 3 in Fig. 1). The reason we design multi-stage semantic GANs is that, we observed that only one semantic GAN layer cannot provide effective semantic supervision. However, existing research [29, 46] also proves that the perceptual loss is helpful for better visual quality on image restoration tasks, nevertheless produces artifacts into the structures and causes color distortion. Hence a proper n is important and will be discussed it in Sec. 4. It is worth noting that semantic constraint has been mentioned in recent work IMAGINE [69], where a semantic constraint is used to guide the category information of syntheses. Whereas in our approach semantic constraint aims to replenish specific semantic details on the layout.

Texture GAN. The last term $\prod_{i=n+1}^N p(y_{tex}^i | x_{tex}^i, y^{i-1})$ indicates that adding texture details to y^n based on the reference x . We follow SinGAN to adopt the the multi-stage architecture with the residual block that is widely used in super-resolution models [22]

$$y^i = G_i(y^{i-1\uparrow}) + y^{i-1\uparrow}, \quad i = n + 1, \dots, N. \quad (6)$$

The training loss contains the adversarial loss and reconstruction loss,

$$\min_{G_i} \max_{D_i} \mathcal{L}_{adv}(G_i, D_i) + \alpha_2 \mathcal{L}_{rec}(G_i). \quad (7)$$

Implementation. To speed up the training, except that D_0 comprises three convolutional layers and one fully connected layer, all other G_i and D_i are simply composed of five convolutional layers. Supposed the height H of x is longer than the width W , the longer side H_N of x^N is set to no more than 256 and the shorter side W_0 of x^0 is fixed to 32. The rescaling method influences both quality of synthesis and training time. See Fig. 3, SinGAN [5] takes




Figure 3: Comparison of different rescaling methods when $N = 5$ and 9 , the horizontal and vertical axes respectively represent stage and pixel.

the basic rescaling method $H_s = H_N \times r^{N-s}$, where r is scalar factor, and s is current stage. ConSinGAN [19] takes a complicated rescaling method to focus on the lower stage. However, we find that these rescaling methods are time-consuming and not robust (see Appendix for details). Hence we design a new rescaling method by analogous Taylor approximation of the basic rescaling method

$$H_s = H_0 \times \left(1 + st + \frac{s(s-1)(s-2)}{k}t^3\right), s = 0, \dots, N. \quad (8)$$

In above equation, $t = \frac{1}{r} - 1$ and k is adjustable to control the curve shape. In this paper we set $k = 2$, r is determined by N . Fig. 3 shows the slope of our rescaling method does not vanish at original point, and becomes larger when s approaches to N , which makes the model more concerned with the generation at intermediate stages and less time-consuming at top stages.

For training structural GAN, we jitter the latent code by Gaussian noise with mean 0, and standard deviations varying from 0.1 to 0.5 to get 500 training data in total. Additional 150 copies of the original images are also injected into training data. Existing works used the convolutional layers of pre-trained VGG-19 [35] for computing perceptual loss. We have tried the different convolutional layers of VGG-19 in advance and concluded that the combination of layers from `relu51` to `relu53` is the best choice for our model. To take trade of training time and performance, we set $N = 6$, $n = 3$, $\lambda = 0.1$, $\alpha_1, \alpha_2 = 10$. For structural GAN, we set training epochs to 10k, learning rate to 10^{-4} , batch size to 32 for both G_0 and D_0 . For the semantic and texture GAN, we set training epochs to 2k, learning rate to 5×10^{-4} for G_i and D_i , $i = 1, \dots, N$.

4 Experimental Results

In this section, we qualitatively and quantitatively compare ExSinGAN with DGP [27], SinGAN [31], ConSinGAN [19], HP-VAEGAN [18] in terms of image synthesis, editing, harmonization and paint-to-image tasks, then conduct ablation study to demonstrate that each module of ExSinGAN is significant.

4.1 Image Synthesis

Qualitative Evaluation. We conduct experiments on various types of images for comprehensive comparison (Fig. 4), including texture image with rich texture but the simple structure, semantic image containing a single complicated object, and structural image comprising a complex scene with several objects.

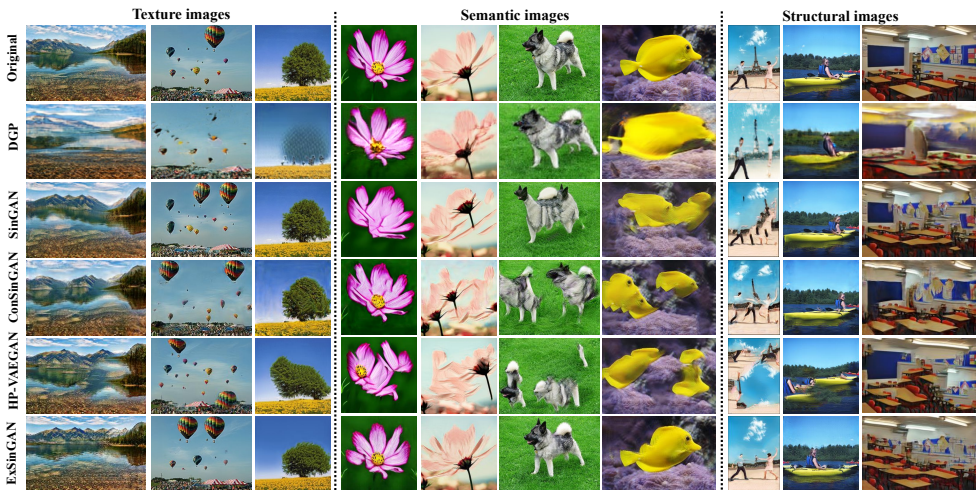



Figure 4: Comparison of DGP [27], SinGAN [51], ConSinGAN [19], HP-VAEGAN [18] and ExSinGAN on image synthesis.



(a). Editing

(b). Harmonization

(c). Paint-to-image

Figure 5: Comparison of SinGAN [51], ConSinGAN [19], ExSinGAN on editing, harmonization and paint-to-image. In (a) and (b), the five columns respectively are original, naive, SinGAN, ConSinGAN, ExSinGAN. In (c), the four columns are original, painting, SinGAN, ExSinGAN.

Firstly, the generation of texture images primarily tests the ability to synthesize texture. Although syntheses of DGP have unpromising texture, ExSinGAN can obtain exquisite results rivaling SinGAN by revising these unsatisfactory images. Note that the syntheses of ConSinGAN have artifacts and coarse texture, and this phenomenon persists when we increase its stages. Secondly, free deformation of the structure and repetitive patches are no longer permitted in semantic images, which further tests the generative ability of models. We first illustrate the erratic generative ability of previous models with two flowers of similar structure in semantic images of Fig. 4. For the first flower, SinGAN and ConSinGAN are confused with the relation of petals and ignore the flower stalk, but the overall layouts of syntheses are plausible. However, they are failed to synthesize the second flower, where the syntheses are of unreasonable repetitive patches petals. In fact, through experiments we find that SinGAN and ConSinGAN are sensitive to color change very much. Note that the color of the first flower is distinct from the background color, hence they can perceive the existence of the petals rather than the stalk. Similarly, they can perceive the existence of the stalk rather than the petals for the second flower. ExSinGAN not only successfully synthesizes the petals but also generates the stalk beneath the petals. For the dog image and fish image, except that

Model	Places50		LSUN50		ImageNet50		Training time
	SIFID↓	SSIM↓	SIFID↓	SSIM↓	SIFID↓	SSIM↓	
DGP[[2]]	0.64	0.18	0.75	0.13	1.65	0.17	2 min
SinGAN[[5]]	0.09	0.37	0.23	0.33	0.60	0.34	80 min
ConSinGAN[[15]]	0.06	0.25	0.11	0.26	0.56	0.23	25 min
HP-VAEGAN[[13]]	0.30	0.19	0.32	0.16	0.60	0.16	63 min
ExSinGAN	0.10	0.34	0.11	0.35	0.45	0.34	40 min
Te.7	0.24	0.19	0.53	0.14	0.51	0.19	-
St.1 Te.6	0.17	0.32	0.17	0.31	0.50	0.27	-
St.1 Se.1 Te.5	0.15	0.34	0.15	0.34	0.54	0.34	-
St.1 Se.3 Te.3	0.10	0.34	0.11	0.35	0.45	0.34	-
St.1 Se.6	0.11	0.35	0.11	0.35	0.50	0.34	-
Te.1 Se.3 Te.3	0.23	0.20	0.28	0.19	0.56	0.20	-

Table 1: Evaluations of *SIFID*, *SSIM* on *Places50* [[\[5\]](#)], *LSUN50* [[\[15\]](#)], and *ImageNet50*. St., Se. and Te. are the abbreviations of Structural, Semantic and Texture GAN. The numbers behind them are their stages.

Model	Places50		LSUN50		ImageNet50	
	Realism↑	Diversity↑	Realism↑	Diversity↑	Realism↑	Diversity↑
SinGAN[[5]]	23.0%	47.8%	18.8%	39.9%	11.5%	35.3%
ExSinGAN	77.0%	52.2%	81.2%	60.1%	88.5%	64.7%

Table 2: User studies on *Places50*, *LSUN50*, and *ImageNet50*. Here the number represents the percentage of votes that voted for corresponding method.

syntheses of DGP have obvious texture defects, the other three generative methods are completely failed in vision. This is similar to case 1 shown in Fig. 1. Due to the lack of structural and semantic constraints, they can only generate repeated patches, while the syntheses of ExSinGAN can also maintain reasonable structure and high-quality texture. Thirdly, for the structural images, despite the fact DGP performs badly in texture, ExSinGAN just properly utilizes its merit of generating plausible layouts, making reasonable variations in images. Visually, ExSinGAN has made satisfying structural changes rather than simply distorting the image. For example, for the dancing image synthesized by ExSinGAN, although the tower in the background disappears, the structure of the image is not polluted like the syntheses of SinGAN. As can be seen from the boat and classroom images, ExSinGAN can produce stronger changes in the layout, proving that ExSinGAN is not overfitting to the given image.

Quantitative Evaluation. Constructing valid quantitative metric for single image generative model is really an open and difficult problem. We adopt the *SIFID* proposed by SinGAN [[\[5\]](#)] to compare the distribution of the original image with that of syntheses using the features extracted by the pre-trained Inception network [[\[6\]](#)]. The classic *SSIM* [[\[10\]](#)] is used for computing the diversity. Models are evaluated on three datasets, *Places50* proposed by SinGAN consisting of 50 images from the *Places* dataset [[\[5\]](#)], *LSUN50* proposed by ConSinGAN consisting 50 images from the *LSUN* dataset [[\[15\]](#)], *ImageNet50* containing 50 semantic images from the validation set of *ImageNet* to test generation of semantic images. The quantitative evaluation in Table 1 shows that ExSinGAN achieves similar *SSIM* to SinGAN but better performance on *SIFID*. In particular, on the most difficult dataset *ImageNet50*, ExSinGAN is ahead of other methods, proving that the hierarchical structure is highly effective to learn a pretty generative model. As previous discussion, DGP performs badly in texture (high *SIFID*) but generates good layouts (low *SSIM*). ExSinGAN takes both the strengths of GAN inversion (better layout) and internal learning methods (better pixel-level generation), performing outstandingly and stably all the time.

User Studies. Moreover, we conduct user studies to compare the realism and diversity of our syntheses with that of SinGAN via Amazon Mechanical Turk. We provide the original image and syntheses from ExSinGAN and SinGAN of the three datasets to Turkers, and Turkers will select the better image with regard to two criteria: realism (the more realistic image) and diversity (the image more different from the given image). Specifically, we recruited 50 Turkers for each dataset, and for each image we provide one synthesis of SinGAN and ExSinGAN respectively. The Turkers will make decisions without the time limitation. The result of user studies in Table 2 shows that our method can synthesize more realistic and diverse images on all datasets. Although the *SSIM* of ExSinGAN is similar to that of SinGAN, it can be attributed to that the syntheses of SinGAN contain artifacts that are meaningless to human vision. Moreover, as the difficulty of datasets increases, the gap between ExSinGAN and SinGAN also broadens, indicating that ExSinGAN has a better generalization performance on various images.

4.2 Generalization Study

Editing. This task aims to produce a seamless composite in which image regions have been copied and pasted in other locations. For ExSinGAN, we inject a subsampled version of the naive image into stage one as x_{str} , then combine the synthesis with the original image at the edited regions. Fig. 5(a) shows some examples of three methods. For the stone image, ExSinGAN performs better on generating more realistic texture than the other two methods. For the tree image, ExSinGAN can keep the structures of naive images well, which means ExSinGAN is more controllable. ConSinGAN failed to merge the background in this task.

Harmonization. Image harmonization aims to realistically blend a pasted object with a background image. For SinGAN and ExSinGAN, the harmonization process is identical to the editing process. The results in Fig. 5(b) show that our work still gets better results compared with SinGAN. ConSinGAN can get better results sometimes because it fine-tunes the models by the naive image.

Paint-to-image. See Fig. 5(c), this task aims to transfer a paint into a photo-realistic image. Because ConSinGAN does not support this task, we just compare ExSinGAN with SinGAN. For SinGAN, this is done by subsampling the paint and feeding it into one of the coarse scales. For ExSinGAN, we just feed the subsampled painting into stage one. From these images, we know that ExSinGAN is sensitive to the structure of paint (the road under trees, the grass near the mountains), and the texture and style are more realistic (the cows). For the pyramids, the synthesis of ExSinGAN is more three-dimensional and context-based (the clouds around pyramids). In conclusion, ExSinGAN can do this task better than SinGAN with fewer stages and time consumption.

4.3 Ablation Study

In this subsection, we conduct both quantitative (Table 1) and qualitative (Fig. 6) ablation study on ExSinGAN to verify that our improvements in SinGAN are effective. We start with the baseline 7-stage SinGAN and check the effects of each modular GAN.

Structural GAN. Fig. 6 shows an awful synthesis of the baseline model. To prove the significance of structural GAN, we just replace the first stage of baseline with structural GAN, then the synthesis (the third image) has an overtly clearer appearance than before, and the *SIFID* decreases overtly (St.1 Te.6 in Table 1).




Figure 6: Examples for ablation study. From left to right, first row: original, SinGAN, structural GAN with texture GAN, ExSinGAN, semantic GAN with texture GAN, using the discriminator of BigGAN for computing the perceptual loss. Second row: original, ExSinGAN with increasing n from 1 to 5. Third row: original, and ExSinGAN with increasing N from 5 to 9.

Semantic GAN. Now we combine the Structural GAN with Semantic GAN to show their effects further. In the first row of Fig. 6, the fourth image (ExSinGAN) shows that the skin of the wolf is smoother and tidier than the previous one. Then we withdraw the structural GAN (Te.1 Se.3 Te.3). The fifth image shows that the result of generation becomes terrible in vision. It means that the structural GAN is necessary for the successful generation, and the semantic GAN is a helper to make the synthesis more realistic and meaningful. They promote and affect each other. In Table 1, the quantitative results show that only Semantic GAN has a negative effect on the baseline, but a positive effect on structural GAN. We also try to use the discriminator of BigGAN to compute the perceptual loss, and the results show that it plays a similar role with VGG-19, meaning it can extract useful semantic information.

Number of Stages. The number of stages greatly influences the results and the cost of time. For ExSinGAN, there are two important hyper-parameters N and n that should be considered. When we fix $N = 6$, the second row of Fig. 6 shows that the semantics of syntheses is closer to the original image as n is increased from 1 to 5. However, since the high-level feature maps also contain information for image restoration, once the size of the feature map is too large, the perceptual loss is almost equivalent to per-pixel loss, which is inconsistent with the purpose of semantic constraint and adds artifacts to the syntheses. The quantitative results (St.1 Te.6, St.1 Se.1 Te.6, St.1 Se.3 Te.3, and St.1 Se.6) show that *SIFID* firstly decreases and then increases, proving that $n = 3$ is the best choice. We do not conduct a quantitative study about N because there has been similar research in [63] and [49]. Here we just support some examples, the third row of Fig. 6 shows that higher-quality syntheses can be obtained with N increased, but of smaller changes of structure and more time consuming (a double-time when $N = 9$).

5 Conclusion

We introduced an explainable single image generative model named ExSinGAN based on the triplet representation of image. ExSinGAN has three modular GANs, respectively structure GAN generating reasonable layouts, semantic GAN replenishing the semantic details, and texture GAN adding the fine texture details. We demonstrated its superior abilities in synthesis and image manipulation tasks compared with previous works. We believe that the combination of external knowledge and internal learning is the key to solving the single image generation task, which will be further explored in our future work.

Acknowledgements. This paper is supported by the National Natural Science Foundation of China (Nos. U19B2040, 11731013, 11991022), the Strategic Priority Research Program of Chinese Academy of Sciences, Grant No. XDA27000000, and the Fundamental Research Funds for the Central Universities.

References

- [1] Connelly Barnes, E. Shechtman, A. Finkelstein, and Dan B. Goldman. Patchmatch: a randomized correspondence algorithm for structural image editing. In *SIGGRAPH 2009*, 2009.
- [2] David Bau, Jun-Yan Zhu, Hendrik Strobelt, B. Zhou, J. Tenenbaum, W. Freeman, and A. Torralba. Gan dissection: Visualizing and understanding generative adversarial networks. *ArXiv*, abs/1811.10597, 2019.
- [3] David Bau, Hendrik Strobelt, William Peebles, Bolei Zhou, Jun-Yan Zhu, Antonio Torralba, et al. Semantic photo manipulation with a generative image prior. *arXiv preprint arXiv:2005.07727*, 2020.
- [4] Sagie Benaim, Ron Mokady, Amit H. Bermano, D. Cohen-Or, and L. Wolf. Structural analogy from a single image pair. *Computer Graphics Forum*, 40, 2021.
- [5] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Neural photo editing with introspective adversarial networks. *arXiv preprint arXiv:1609.07093*, 2016.
- [6] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=B1xsqj09Fm>.
- [7] Leilei Cao, Tong Yang, Yixu Wang, Bo Yan, and Yandong Guo. Generator pyramid for high-resolution image inpainting. *ArXiv*, abs/2012.02381, 2020.
- [8] J. Chen, Qihui Xu, Qi Kang, and M. Zhou. Mogan: Morphologic-structure-aware generative learning from a single image. *ArXiv*, abs/2103.02997, 2021.
- [9] Antonia Creswell and Anil Anthony Bharath. Inverting the generator of a generative adversarial network. *IEEE transactions on neural networks and learning systems*, 30(7):1967–1974, 2018.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [11] Emily L. Denton, Soumith Chintala, Arthur D. Szlam, and R. Fergus. Deep generative image models using a laplacian pyramid of adversarial networks. In *NIPS*, 2015.
- [12] J. Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019.
- [13] J. Donahue and K. Simonyan. Large scale adversarial representation learning. In *NeurIPS*, 2019.

- [14] Jeff Donahue, Philipp Krähenbühl, and Trevor Darrell. Adversarial feature learning. *arXiv preprint arXiv:1605.09782*, 2016.
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [16] Niv Granot, Assaf Shocher, Ben Feinsein, S. Bagon, and M. Irani. Drop the gan: In defense of patches nearest neighbors as single image generative models. *ArXiv*, abs/2103.15545, 2021.
- [17] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777, 2017.
- [18] Shir Gur, Sagie Benaim, and L. Wolf. Hierarchical patch vae-gan: Generating diverse videos from a single sample. *ArXiv*, abs/2006.12226, 2020.
- [19] Tobias Hinz, Matthew Fisher, Oliver Wang, and Stefan Wermt. Improved techniques for training single-image gans. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1300–1309, 2021.
- [20] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [21] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- [22] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1646–1654, 2016.
- [23] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [24] Zachary C Lipton and Subarna Tripathi. Precise recovery of latent vectors from generative adversarial networks. *arXiv preprint arXiv:1702.04782*, 2017.
- [25] Sutharsan Mahendren, C. Edusooriya, and R. Rodrigo. Diverse single image generation with controllable global structure through self-attention. *ArXiv*, abs/2102.04780, 2021.
- [26] J. O’Regan and A. Noë. A sensorimotor account of vision and visual consciousness. *The Behavioral and brain sciences*, 24 5:939–73; discussion 973–1031, 2001.
- [27] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo. Exploiting deep generative prior for versatile image restoration and manipulation. *arXiv preprint arXiv:2003.13659*, 2020.

- [28] Ankit Raj, Yuqi Li, and Yoram Bresler. Gan-based projector for faster recovery with convergence guarantees in linear inverse problems. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5602–5611, 2019.
- [29] Mehdi S. M. Sajjadi, B. Schölkopf, and M. Hirsch. Enhancenet: Single image super-resolution through automated texture synthesis. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 4501–4510, 2017.
- [30] Viraj Shah and Chinmay Hegde. Solving linear inverse problems using gan priors: An algorithm with provable guarantees. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 4609–4613. IEEE, 2018.
- [31] Tamar Rott Shaham, Tali Dekel, and T. Michaeli. Singan: Learning a generative model from a single natural image. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4569–4579, 2019.
- [32] Assaf Shocher, N. Cohen, and M. Irani. "zero-shot" super-resolution using deep internal learning. In *CVPR*, 2018.
- [33] Assaf Shocher, Nadav Cohen, and Michal Irani. "zero-shot" super-resolution using deep internal learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3118–3126, 2018.
- [34] Assaf Shocher, S. Bagon, Phillip Isola, and M. Irani. Ingan: Capturing and retargeting the "dna" of a natural image. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4491–4500, 2019.
- [35] K. Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2015.
- [36] V. Sushko, Juergen Gall, and A. Khoreva. One-shot gan: Learning to generate samples from single images and videos. *ArXiv*, abs/2103.13389, 2021.
- [37] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [38] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9446–9454, 2018.
- [39] Pei Wang, Yijun Li, Krishna Kumar Singh, Jingwan Lu, and N. Vasconcelos. Imagine: Image synthesis by image-guided model inversion. *ArXiv*, abs/2104.05895, 2021.
- [40] Zhou Wang, A. Bovik, H. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13:600–612, 2004.
- [41] Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. Gan inversion: A survey. *arXiv preprint arXiv:2101.05278*, 2021.

- [42] Rui Xu, Xintao Wang, Kai Chen, B. Zhou, and Chen Change Loy. Positional encoding as spatial inductive bias in gans. *ArXiv*, abs/2012.05217, 2020.
- [43] Raymond A Yeh, Chen Chen, Teck Yian Lim, Alexander G Schwing, Mark Hasegawa-Johnson, and Minh N Do. Semantic image inpainting with deep generative models. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5485–5493, 2017.
- [44] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.
- [45] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning deep features for scene recognition using places database. *Advances in neural information processing systems*, 27:487–495, 2014.
- [46] Yang Zhou, Zhen Zhu, Xiang Bai, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Non-stationary texture synthesis by adversarial expansion. *ACM Transactions on Graphics (TOG)*, 37(4):1–13, 2018.
- [47] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A Efros. Generative visual manipulation on the natural image manifold. In *European conference on computer vision*, pages 597–613. Springer, 2016.