SCIENCE CHINA Information Sciences



• REVIEW •

May 2023, Vol. 66 151101:1–151101:28 https://doi.org/10.1007/s11432-022-3679-0

Survey on leveraging pre-trained generative adversarial networks for image editing and restoration

Ming LIU^{1,2}, Yuxiang WEI¹, Xiaohe WU¹, Wangmeng ZUO^{1*} & Lei ZHANG²

¹School of Computer Science and Technology, Harbin Institute of Technology, Harbin 150001, China; ²Department of Computing, Hong Kong Polytechnic University, Hong Kong 999077, China

Received 29 July 2022/Revised 25 October 2022/Accepted 7 December 2022/Published online 4 April 2023

Abstract Generative adversarial networks (GANs) have drawn enormous attention due to their simple yet effective training mechanism and superior image generation quality. With the ability to generate photo-realistic high-resolution (e.g., 1024×1024) images, recent GAN models have greatly narrowed the gaps between the generated images and the real ones. Therefore, many recent studies show emerging interest to take advantage of pre-trained GAN models by exploiting the well-disentangled latent space and the learned GAN priors. In this study, we briefly review recent progress on leveraging pre-trained large-scale GAN models from three aspects, i.e., (1) the training of large-scale generative adversarial networks, (2) exploring and understanding the pre-trained GAN models, and (3) leveraging these models for subsequent tasks like image restoration and editing.

Keywords survey, generative adversarial networks, pre-trained models, image editing, image restoration

Citation Liu M, Wei Y X, Wu X H, et al. Survey on leveraging pre-trained generative adversarial networks for image editing and restoration. Sci China Inf Sci, 2023, 66(5): 151101, https://doi.org/10.1007/s11432-022-3679-0

1 Introduction

The recent years have witnessed the rapid development of generative adversarial networks (GANs) [1]. Particularly, GAN models are trained by learning to push a fake data distribution towards the real one via adversarial learning, thereby making the generated data (e.g., images) indistinguishable from the real one. In general, the GAN models are composed of two parts, i.e., a generator G and a discriminator D. Specifically, the generator aims at generating fake samples as realistically as possible, while the discriminator, in turn, aims at distinguishing between real and fake samples to avoid being fooled by the generator. During the adversary process, the ability of G and D can be constantly improved. Finally, they may reach a balanced state (i.e., Nash equilibrium), where G is capable of generating images that D cannot tell between real and fake. A series of studies made tremendous efforts to improve the quality of GAN-generated images in terms of network structures [2–5], loss functions [6–11], and so on. Recently, several GAN models (e.g., PGGAN [12], BigGAN [13], and StyleGAN series [14–17]) have shown superior image generation ability to produce photo-realistic high-resolution (e.g., 1024×1024) images.

Therefore, the exploration of leveraging GANs for enhancing image quality has drawn upsurging attention. Prior studies often simply take the adversarial loss as an extra term in their learning objectives. On the one hand, the discriminator and adversarial loss can serve as out-of-the-box components, which could be readily incorporated into many other models, and bring no extra complexity during inference. On the other hand, applying adversarial training generally can avoid the models being optimized to a blurry average solution, which makes the output image contain more sharp and clear details. As such, the effectiveness has been verified in many vision tasks such as image editing [18–23], image superresolution [24–27], image deblurring [28–30], and image dehazing [31–34].

^{*} Corresponding author (email: wmzuo@hit.edu.cn)



Figure 1 (Color online) Image interpolation in the pixel space and the latent space. x_1 and x_2 are two images of the same object. E(x) denotes the latent code of an image x, which is mapped backed to the latent space by the encoder E. G(E(x)) denotes the inversion result of x. The first row shows the results of interpolating directly in the pixel space via $x = \alpha \cdot x_1 + (1 - \alpha) \cdot x_2$. In contrast, the second row shows the results of interpolating in the latent space via $x = G(\beta \cdot E(x_1) + (1 - \beta) \cdot E(x_2))$. It can be seen that the interpolation in the latent space is semantically more plausible. The sample images are from the tower subset of LSUN dataset [37], and the results in the second row are generated via the officially released model of IDInvert [38].

However, with the development of GANs, it becomes much harder for such a utilization scheme to fully explore the potential of GAN models. Therefore, many recent studies turn to exploring and leveraging pre-trained GAN models, whose advantages are analyzed as follows.

Disentanglement of GAN latent space. For many vision tasks like image editing, it is a basic problem to identify image contents and the manipulation mode of a concept to be edited. Since the concepts are entangled with each other and form an interlaced group of editing spaces (e.g., changing a facial image to smile or old can both cause wrinkles). Even with the well-labeled datasets (e.g., CelebA [35]), it is still a very challenging task to describe the manipulation mode directly on the pixels [20–23]. As a remedy, pre-trained GAN models provide yet another possible solution, where the disentanglement of latent spaces has been explored recently [14, 36]. As shown in Figure 1 [37, 38], interpolation between images will cause obvious artifacts, while the same operation in the latent space will be semantically more plausible. Such a phenomenon clearly illustrates the latent space disentanglement ability of recent GAN models.

Leveraging pre-trained GAN priors. For tasks with image outputs, regularization terms and natural image priors are often applied to enhance the output image quality, e.g., local smoothness [39], non-local self-similarity [40], sparsity [41], etc. Nevertheless, they are either hand-crafted with manually tuned parameters, or based on strong assumptions that are usually not applicable in real applications, making the performance and flexibility greatly limited. Some unsupervised methods [42] propose to get rid of these limitations, but they are designed and evaluated on specific degradations. Dynamic network [43] or network architecture search (NAS) [44] based methods introduce structure flexibility, but they are still limited in generalizing to a large variety of degradations. On the contrary, the learning objective of GAN models is to generate high-quality diverse images from the latent code (which is generally a random noise). In other words, given a vector located in the latent space, the well-trained GAN models are very likely to generate a high-quality output. Therefore, the pre-trained GAN models are naturally able to provide stronger and more reliable high-quality natural image priors. Furthermore, the GAN priors can be helpful to a wide range of tasks, making it with much more practical significance.

Training scheme & development potentials. Collecting pairwise training data is extremely timeconsuming and expensive, and it is impossible to collect abundant training data for all tasks. Fortunately, GAN models are trained in an unsupervised manner, where the training images are much easier to collect without requiring ground-truth labels or reference images. Indeed, there are already a huge amount of unlabeled images, and the unsupervised training scheme of GAN models makes it easier to apply the extremely large-scale training sets and to train models with greater capacity for better performance.

As a result, many recent studies show rising interest in pre-trained GAN models and have developed a series of methods. In this study, we review the recent progress on leveraging pre-trained GAN models, from understanding GAN models [45–48] to leveraging them for subsequent tasks such as image editing [49–52] and restoration [53–57], mainly involving,

Liu M, et al. Sci China Inf Sci May 2023 Vol. 66 151101:3

Format	Definition	Example			
Upper-case italic letters (e.g., G) Models or networks	Generator G ; discriminator D ; encoder E ;			
		MLP F ; segmentation model S			
Bold letters (e.g., \boldsymbol{z})	Images or tensors	Image $\boldsymbol{x}, \boldsymbol{I}$; latent code (or random noise) $\boldsymbol{z}, \boldsymbol{w}, \boldsymbol{p}, \boldsymbol{n},$			
		etc.; learned input m ; Fourier feature input F ; condition c			
Handwritten letter (e.g., \mathcal{Z}) Latent spaces or distribution		Latent spaces $\mathcal{Z}, \mathcal{C}, \mathcal{W}, \mathcal{S},$			
		$\mathcal{P}, \mathcal{N}, \mathcal{F};$ loss functions \mathcal{L}			
Greek letters (e.g., α) Parameters or spec		Interpolation parameters or hyperparameters α , β ,			
		$\lambda; \ oldsymbol{ heta}$ model parameters; any possible network inputs $oldsymbol{\delta}$			
Double struck letter (e.g., $\mathbb{R})$	Following mathematical definitions	Mathematical expectation \mathbb{E} ; spatial dimension \mathbb{R}			

Table 1	A summarv	on the	symbols	used in	this survey
raoro r	II DUILINGUL y	011 0110	0,110010	about III	. orres boar tog

• We survey the characteristics of pre-trained GAN models, and the efforts for exploring and acquiring these characteristics are also introduced.

• Focusing on image restoration and editing tasks, we summarize and compare relevant methods to utilize pre-trained GAN models.

• We discuss the open problems and potential research directions in this field.

The rest of this study is organized as follows. To begin with, Section 2 introduces recent progress of GAN models, including network architecture, and training mechanism. Then, the latent space of GANs as well as methods for GAN inversion are introduced in Section 3. By exploring the characteristics of pre-trained GAN models (e.g., latent space exploration), the relevant methods and experiments for understanding pre-trained GAN models are reviewed in Section 4. Based on the impressions of latent space, we introduce and summarize the methods in Section 5, which leverage pre-trained GANs for subsequent applications, especially the image editing and restoration tasks. Finally, we discuss the open problems in this field, and point out some potential directions for future research in Section 6. More information about relevant methods and repositories can be found at the website ¹).

2 Pre-trained generative adversarial networks

2.1 Preliminary

In the deep learning [58] era, when researchers are referring to a task, generally there accompany two fundamental problems, i.e., data sets and evaluation metrics. For GAN models, these two problems are even more important. On the one hand, numerous training samples are typically required to avoid overfitting and to tap the potential of their superior capacity. On the other hand, the evaluation of generative models is more difficult since there are no ground-truth or labels for assessment, and generating thousands or even millions of images for user study is a mission impossible due to the time and labor costs. Therefore, before introducing the large-scale GAN models, we brief the relevant datasets, as well as the metrics for evaluating the pre-trained GAN models. Besides, we summarize the symbols used in this survey in Table 1.

2.1.1 Image datasets for training large-scale GAN models

ImageNet [59] is the very first large-scale natural image dataset, which contains over 14M images from more than 21k classes (referred to as ImageNet-21k), and the average image size is 469×387 . The commonly used configuration is a 1000-class subset with ~1k images per class (denoted by ImageNet-1k), and the images are usually resized to 224×224 or 256×256 in practice.

CelebA [35] (large-scale CelebFaces Attributes dataset) contains more than 200k facial images from 10k celebrities, each of the images is annotated with 40 attributes, and the aligned images are resized and cropped to 178×218 . A subset of CelebA is post-processed to form the CelebA-HQ [12] dataset, containing 30k facial images with the resolution of 1024×1024 . The dataset is also extended with annotations like CelebAMask-HQ [60] via pixel-wise facial component labeling (face parsing).

LSUN [37] & AFHQ [23] are another two datasets of scene and object categories. The LSUN dataset is composed of 10 scene categories and 20 object categories, each of which contains around 1M labeled

¹⁾ https://github.com/csmliu/pretrained-GANs.

images, and the images are provided by resizing the shorter edge to 256 pixels and compressing to JPEG image quality of 75. On the contrary, AFHQ (Animal Faces-HQ dataset) is composed of only three categories (i.e., cat, dog, and wildlife), each of which contains around 5k high-quality animal face images with a resolution of 512×512 . Recent methods often take a specific category of these two datasets for training their GAN models, e.g., bedroom, church from LSUN, and cat from AFHQ.

FFHQ [14] (Flickr-Faces-HQ dataset) is a high-quality image dataset collected from Flickr²⁾, containing 70k automatically aligned high-resolution (1024×1024) facial images with diverse data distribution. To date, FFHQ is the most popular dataset for training GAN models for facial image generation. Image editing and restoration methods based on pre-trained GANs are often trained on FFHQ dataset and evaluated on CelebA-HQ [12] to show the generalization ability.

Other datasets. Apart from the aforementioned ones, some other datasets [16, 61–65] are also employed for prototyping and/or training large-scale GAN models. To sum up, large scale and favorable image quality are two key ingredients of datasets for training high-quality large-scale GAN models. Therefore, some other datasets like Objects 365 [66], Places 365 [67], and Open Images [68] may also be used for training GAN models with greater capacity.

2.1.2 Evaluation metrics for assessing GAN models

Mean opinion score (MOS) is the most intuitive index to assess the perceptual quality of images, since the image quality assessment is conducted by human raters. However, assessing via MOS is very expensive, and the results may be biased due to factors like subjective perceiving differences. Furthermore, the time-consuming procedure makes it suitable for small-scale assessment, e.g., user study, but hard to be employed for the evaluation during training and broader comparison.

Inception Score [69] (IS) is a commonly used metric for evaluating the image quality and class diversity. Considering a well-trained classifier (e.g., Inception-v3 [70] trained on ImageNet [59] for calculating IS), the high confidence of classifying a generated image \hat{x} into a class c (i.e., smaller entropy of $p(c|\hat{x})$) generally implies decent image generation quality. On the other hand, better class diversity requires that the marginal probability distribution of class label approximates uniform distribution, i.e., higher entropy of p(c). Therefore, the Inception Score is defined by

$$IS = \exp(\mathbb{E}_{\hat{\boldsymbol{x}} \sim p_q}[D_{KL}(p(c|\hat{\boldsymbol{x}}) \| p(c))]) = \exp(H(c) - \mathbb{E}_{\hat{\boldsymbol{x}} \sim p_q}[H(c|\hat{\boldsymbol{x}})]),$$
(1)

where \hat{x} is sampled from the generated image distribution p_g , $D_{\text{KL}}(\cdot \| \cdot)$ denotes Kullback-Leibler (K-L) divergence, $H(\cdot)$ represents entropy. However, IS is unable to detect mode collapse, e.g., a high Inception Score will be derived when the generator produces one and only one high-quality image for each class.

Modified inception score [71] (m-IS) modifies IS by introducing the cross-entropy style score, i.e., $-p(c|\mathbf{x}_i)\log(p(c|\mathbf{x}_j))$, which considers the diversity of a particular class. Formally, m-IS is defined by

$$m-IS = \exp(\mathbb{E}_{\hat{\boldsymbol{x}}_i \sim p_q}[\mathbb{E}_{\hat{\boldsymbol{x}}_i \sim p_q}[D_{KL}(p(c|\hat{\boldsymbol{x}}_i) \| p(c|\hat{\boldsymbol{x}}_j))]]).$$
(2)

Note that Eq. (2) calculates the m-IS on a particular class, and the final m-IS is obtained by averaging the scores on all classes.

Mode Score [72] (MS) proposes to improve IS by introducing the prior distribution of the real images,

$$MS = \exp(\mathbb{E}_{\hat{\boldsymbol{x}} \sim p_g}[D_{KL}(p(c|\hat{\boldsymbol{x}}) \| p(c^{\text{train}}))] - D_{KL}(p(c) \| p(c^{\text{train}}))),$$
(3)

where p(c) is the same as IS, i.e., the distribution obtained from generated samples, while $p(c^{\text{train}})$ is the distribution of the training data. In this way, MS can evaluate both variety and visual quality in a single metric. However, IS and MS were proven equivalent by Zhou et al. [73], and they further proposed AM Score in [74].

AM Score [74] replaces the entropy H(c) in (1) with the K-L divergence between c and c^{train} , which is defined by

$$AM = D_{KL}(p(c^{\text{train}}) \| p(c)) + \mathbb{E}_{\hat{\boldsymbol{x}} \sim p_g}[H(c | \hat{\boldsymbol{x}})], \qquad (4)$$

where the first term requires that c^{train} is close to c, while the second term requires that the class label predicted from \hat{x} has low entropy. Besides, AM Score also shows that the entropy should be obtained by a classifier trained on the given dataset, rather than a general classifier (e.g., trained on ImageNet).

²⁾ https://www.flickr.com/.

Fréchet inception distance [75]. Considering the drawbacks of IS, Fréchet inception distance (FID) takes the real images into consideration and calculates the distance between generated images and real ones. Specifically, FID is defined by

$$\operatorname{FID} = \|\mu_r - \mu_g\|_2^2 + \operatorname{Tr}(\Sigma_r + \Sigma_g - 2(\Sigma_r \Sigma_g)^{\frac{1}{2}}), \tag{5}$$

where $\text{Tr}(\cdot)$ denotes trace of a matrix. The real and generated image features extracted by Inceptionv3 [70] are both assumed to follow the multivariate Gaussian distribution, i.e., $\phi(\boldsymbol{x}) \sim \mathcal{N}(\mu_r, \Sigma_r), \phi(\hat{\boldsymbol{x}}) \sim \mathcal{N}(\mu_g, \Sigma_g)$. By utilizing the real samples, FID is not restricted by the training dataset (i.e., ImageNet [59] for Inception-v3 [70]), and is much more sensitive to mode collapse.

Sliced Wasserstein distance [76]. It is worth noting that the Gaussian assumption of FID does not always hold true. To avoid such an assumption, Bonneel et al. [76] proposed to calculate the sliced Wasserstein distance (SWD) to approximate the Wasserstein distance, i.e.,

$$SWD(\boldsymbol{I}_x, \boldsymbol{I}_y) = \left(\int_{\mathbb{S}^{d-1}} W_p^p(\mathcal{R}\boldsymbol{I}_x(\cdot, \theta), \mathcal{R}\boldsymbol{I}_y(\cdot, \theta)) \mathrm{d}\theta \right)^{\frac{1}{p}},$$
(6)

where \mathbb{S}^{d-1} is the unit sphere in \mathbb{R}^d , p is set to 2 in practice, \mathcal{R} denotes the Radon transform. Please refer to [76,77] for more details.

GAN-train & GAN-test [78]. Another problem in GAN training is over-fitting. In other words, the model memorizes the training samples for generation. The GAN-train and GAN-test indices provide a more comprehensive evaluation on the GAN model. In particular, GAN-train is the accuracy of a classifier trained on a generated dataset \mathcal{D}_g and evaluated on a real-image validation dataset $\mathcal{D}_r^{\text{val}}$, while GAN-test is the accuracy of a classifier trained on real-image training set $\mathcal{D}_r^{\text{train}}$ and tested on \mathcal{D}_g . These two metrics can identify problems such as mode dropping, poor image quality, and over-fitting.

Fréchet segmentation distance [46]. As shown in [46], the mode collapse occurs not only at the distribution level, but also at the instance level. In particular, certain objects may not be generated by a GAN model, and this phenomenon is termed by mode dropping. For assessing the mode dropping in a GAN model, the Fréchet segmentation distance (FSD) is defined by modifying FID, i.e.,

$$FSD = \|\mu_g - \mu_t\|_2^2 + Tr(\Sigma_g + \Sigma_t - 2(\Sigma_g \Sigma_t)^{\frac{1}{2}}),$$
(7)

where μ_t (μ_g) is the mean pixel count for each object class over a sample of training images (generated images), and Σ_t (Σ_g) is the covariance of these pixels.

Besides, many other metrics like perceptual path length [14], linear separability [14], precision and recall [79], and geometry score [80] are also explored for better evaluating GANs. For certain tasks like image restoration, peak signal-to-noise ratio (PSNR), structural similarity (SSIM) [81], and learned perceptual image patch similarity (LPIPS) [82] are also employed for fidelity evaluation. We recommend [83] for a more comprehensive survey on GAN metrics.

2.2 Large-scale GAN models

GANs are originally proposed for image generation. As such, we take image generation as an example to introduce the basic concepts and training schemes of GANs, which could be naturally generalized to other domains. A typical GAN model [1,3] is composed of a generator and a discriminator, which are denoted by G and D, respectively. Given a random vector $z \in \mathbb{Z}$, the generator maps it to an image \hat{x} ,

$$\hat{\boldsymbol{x}} = G(\boldsymbol{z}, \boldsymbol{\delta}; \boldsymbol{\theta}_G), \tag{8}$$

where $\boldsymbol{\delta}$ denotes other possible inputs (e.g., class condition, layer-wise noise), $\boldsymbol{\theta}_G$ represents the parameters of G, $\hat{\boldsymbol{x}} \in \mathbb{R}^{H \times W \times C}$ denotes the generated C-channel image with the size of $H \times W$. Then, the discriminator D takes $\hat{\boldsymbol{x}}$ (or a real image \boldsymbol{x}) as input, and predicts the probability that the input is from the distribution of real images or the distance to the real image distribution³⁾, and the learning objective is defined by

$$\min_{G} \max_{D} \mathcal{L}_{\text{GAN}}(G, D) = \mathbb{E}_{\boldsymbol{x} \sim p_{\text{data}}(\boldsymbol{x})}[\log D(\boldsymbol{x})] + \mathbb{E}_{\boldsymbol{z} \sim p_{z}(\boldsymbol{z})}[\log(1 - D(G(\boldsymbol{z})))].$$
(9)

³⁾ Precisely speaking, the metric is diversity rather than distance for some adversarial loss functions such as vanilla GAN [1].



Figure 2 (Color online) Illustration of recent GAN models (see (a)–(d)) and the latent spaces of StyleGAN series [14] (see (e)). (a) For PGGAN [12], the blue part denotes the progressive growing procedure from 4×4 to 8×8 . The components with dash lines are employed for the fade-in strategy, where α is gradually growing to 1. They are discarded when the model grows to a higher resolution. (b) For BigGAN [13], a specific noise is delivered to each layer together with the class embedding, and the model is end-to-end trained without the progressive growing procedure. (c) For StyleGAN [14], a series of FC layers are deployed to map z into w. The green part only belongs to StyleGAN2. (d) For StyleGAN3 [17], the generator is largely modulated to improve the translational and rotation equivariance. The discriminator is omitted since it is identical with that used in StyleGAN2 [15]. (e) For simplicity, here we take the StyleGAN series as an example to show the latent spaces based on the GAN inversion task.

In the following, we mainly introduce a handful of recent GAN models that are used in the methods surveyed in Sections 3–5. We recommend [84,85] for a comprehensive review of GANs.

PGGAN [12]. Leveraging the progressive growing strategy, PGGAN (or ProGAN) achieves the 1024×1024 image generation resolution. Specifically, at the beginning of the training procedure, a 4×4 "image" is generated by the initial generator $G^{(4)}$, and the discriminator $D^{(4)}$ accordingly takes 4×4 input as well. Once the model converges at a resolution, another layer is deployed at the end of the generator, which generates images with $2 \times$ resolution. Finally we can obtain $G^{(1024)}$, which could generate high-quality 1024×1024 images. The overall architecture is shown in Figure 2(a).

Table 2 Comparison of recent GAN models. The two numbers of parameters for StyleGAN3 [17] are respectively for the translational equivariant configuration and the rotation equivariant configuration. z, m, and F denote random noise, learned constant, and Fourier feature input [88], respectively.

GAN model	Publication	Resolution	#Params (M)	Training scheme	Input	Condition	Regularization
PGGAN [12]	ICLR 2018	1024	23.1	Growing	z	_	Grad. Penalty (GP)
BigGAN [13]	ICLR 2019	512	82.5	End-to-end	z	$oldsymbol{c}+oldsymbol{z}$	Ortho. Reg. & R_1
StyleGAN $[14]$	CVPR 2019	1024	26.2	Growing	m	w+n	Mixing Reg., GP & R_1
StyleGAN2 $[15]$	$\mathrm{CVPR}\ 2020$	1024	30.0	Multi-scale	m	w+n	Mixing Reg., Lazy R_1 & PPL
StyleGAN3 $[17]$	NeurIPS 2021	1024	22.3/15.8	End-to-end	F	w	R_1

BigGAN [13]. As the name implies, BigGAN analyzes the influence of GAN model size on generation quality and proposes an architecture with more parameters and end-to-end trained with a large batch size. For more stable training, the top singular value of parameters is used to monitor the mode collapse during training, and a novel regularization term on the generator G is designed according to the analysis. The R_1 zero-centered gradient penalty [86] is also leveraged in the discriminator D. Besides, to achieve class-conditioned generation, BigGAN delivers the class embedding and noise vector into each residual block of the generator. Finally, with the large capacity, BigGAN could generate appealing images with the resolution of 512×512 . The overall architecture is shown in Figure 2(b).

StyleGAN [14]. The StyleGAN series [14,15,17] are the most popular GAN models in recent years. Instead of taking noise as input, Karras et al. [14] proposed to learn a constant input m (as shown in Figure 2(c)). And they used a mapping network composed of several fully-connected layers to map the noise z into a latent representation w, which is delivered to the AdaIN [87] layer in each scale. In addition, layer-wise per-pixel noise was introduced in each scale for further performance improvement and achieved stochastic variation on generated details. StyleGAN follows the progressive growing training scheme of PGGAN [12], and the final-stage structure of StyleGAN is shown as the black part in Figure 2(c).

StyleGAN2 [15]. The overall structure of StyleGAN2 is similar to StyleGAN, and the block structure and regularization terms are modulated for better generalization quality. Besides, StyleGAN2 deploys a "toRGB" module in each scale for substituting the progressive growing strategy. The lazy regularization and perceptual path length regularization are also deployed. Furthermore, Karras et al. [16] proposed StyleGAN2-Ada based on a delicately designed discriminator augmentation mechanism, which largely improves the generated image quality and significantly stabilizes training in limited data regimes.

StyleGAN3 [17]. Since the StyleGAN architecture was proposed, utilizing pre-trained GAN models for downstream tasks has drawn much attention, making the equivariance a necessity for many applications. However, the synthesis process of previous GAN models depends on the coordinates of pixels rather than the surface of the generated objects. Karras et al. [17] analyzed the cause of such conditions, and designed a model with translational equivariance and rotation equivariance, which achieves comparable performance against StyleGAN2 but with even less parameters.

The comparison between these GAN models is shown in Figure 2 and Table 2 [88]. It can be seen that, in order to achieve realistic generation, these methods keep exploring smoother and more stable information flow pipelines and architectures, more disentangled and controllable input and condition forms, as well as better training schemes. In Sections 3 and 4, we will delve more deeply into the GAN models, trying to analyze and understand the characteristics and properties of pre-trained GANs.

3 GAN inversion

As introduced in Section 1, a natural choice for achieving better output image quality in recent years is to apply adversarial losses, especially the image editing and restoration methods [18, 19, 24, 28]. However, such a method is insufficient to fully explore the image generation ability of GAN models. Since the working scheme of GAN models is to map the random vectors (latent representations) into images (Eq. (8)), the very first thing to utilize pre-trained GAN models is to invert the input images back into a meaningful latent space (i.e., GAN inversion). Then the latent representations could be manipulated or optimized for achieving tasks like image editing or restoration. Thus we first introduce the commonly used latent spaces in the literature, and then introduce the typical GAN inversion methods.

3.1 Latent spaces

In order to invert an image back to the latent representation, the first thing is to determine the inversion target, i.e., the latent space. For this purpose, the reconstruction accuracy, interpretability, as well as editability should be considered. In the following, we introduce the commonly used latent spaces in the literature, and an illustration can be found in Figure 2(e).

 \mathcal{Z} space & \mathcal{C} space. Following the vanilla GAN [1], early GAN models [2–5, 12, 13, 89] take random noise vectors as the input, and generate fake images via a stack of convolution layers, i.e.,

$$\hat{\boldsymbol{x}} = G(\boldsymbol{z}, \boldsymbol{c}; \boldsymbol{\theta}_G), \tag{10}$$

where c is the condition information (e.g., class label, attribute annotations) for conditional GANs [13,89]. Therefore, a natural choice is to directly invert the generation process, and map the image back to a random noise z, which spans the Z space. Since z is sampled from a very simple distribution (e.g., the standard Gaussian distribution), the semantic features are largely entangled in the Z space, and the simple Z space is too limited to simultaneously represent both content and semantic information. An alternative way is to jointly use the Z space and the C space, where $c \in C$. Thus some methods (e.g., IcGAN [90]) could invert the image into both Z and C spaces, showing extraordinary performance in disentangling the content and attribute representations. However, these methods require massive human efforts in labeling the attributes to model the C space.

 \mathcal{W} space & $\mathcal{W}+$ space. To separate the semantic conditions (e.g., class, style) with image contents (e.g., identity), Karras et al. [14, 15, 17] proposed the StyleGAN series, which could be formulated by

$$\hat{\boldsymbol{x}} = G(\boldsymbol{m}, \boldsymbol{n}, F(\boldsymbol{z}); \boldsymbol{\theta}_G) = G(\boldsymbol{m}, \boldsymbol{n}, \boldsymbol{w}; \boldsymbol{\theta}_G),$$
(11)

where \boldsymbol{m} and \boldsymbol{n} are learnable constant features and layer-wise noises⁴), F is the multi-layer perceptron (MLP) for mapping \boldsymbol{z} to \boldsymbol{w} . Due to the superior generation quality of StyleGAN models, they become the most popular GAN architectures. Instead of mapping back to the \boldsymbol{Z} space, the latent representation $\boldsymbol{w} \in \boldsymbol{W}$ is generally utilized for StyleGAN inversion, which is proven semantically better disentangled with the affine transmissions by the mapping network F [14]. Based on the \boldsymbol{W} space, some researchers [91] proposed to predict an individual latent representation for each generator layer (see Figure 2(c)), resulting in the \mathcal{W} + space, which shows a better inversion accuracy comparing to the \boldsymbol{W} space. Note that the latent codes in the earlier layers (i.e., the ones near the input end) control coarser-grained attributes, while the finer-grained attributes are controlled by the latent codes in the latter layers (i.e., the ones near the output end), and we refer to [14] for detailed illustrations.

S space & \mathcal{P} space. Although \mathcal{W} and \mathcal{W} + spaces have better properties, the changes in w tend to influence the whole image, and different dimensions in w follow various distributions. As introduced in Subsection 4.1, Bau et al. have explored the relationship between neurons and objects in the generated images. Inspired by the discovery that each neuron (channel) may be related to a specific class or semantic component, several studies choose to predict channel-wise style parameters [92–94] denoted by S space, with which one can control the local content of the images. Besides, the S space is also extended by considering both channel-wise and spatial-wise diversity [95], which is defined by the $S\mathcal{A}$ space, but we still regard it as a kind of S space. Zhu et al. [96] focused on the distribution of the \mathcal{W} space, and proposed to use the latent representation (denoted by $p \in \mathcal{P}$) before the last leaky ReLU layer of the mapping network F. Compared to \mathcal{W} , \mathcal{P} has a simpler structure (i.e., similar to a multivariate Gaussian distribution). And they further propose a \mathcal{P}_N space by mapping the distribution of each latent representation to be of zero mean and unit variance. The \mathcal{P}_N space is better filled by the latent codes, and the image generation quality could be evaluated by the distance between the latent representation and the average latent code.

 \mathcal{N} space & \mathcal{F} space. For current GAN models, the largest dimension of the latent representation is 18 × 512 (the 512-dim w for 18 layers, and the dimension of z is much smaller). Even though the most prominent features of the image have been reconstructed, the high-frequency features cannot be faithfully reproduced by such latent spaces. Therefore, some studies [57,97] proposed to leverage the layer-wise noise maps in StyleGAN and StyleGAN2 (i.e., $n \in \mathcal{N}$), and achieved much more accurate reconstructions. Yet the \mathcal{N} space weakens the editability to some extent; thus another category considers the features directly. For example, Bau et al. [46] inverted the images into the features at a particular

⁴⁾ For StyleGAN3 [17], *m* denotes the Fourier feature [88], while *n* is deprecated.

layer to investigate the mode-dropping problem. Kang et al. [98] proposed to map the image back to the feature map f at a certain scale (e.g., 8×8), and jointly use the \mathcal{F} and \mathcal{W} + spaces (denoted by \mathcal{F}/\mathcal{W} + space).

Some methods also consider fine-tuning part of [50, 99] or the whole [57, 100] GAN model, and let Θ space denote the model parameters. As analyzed in Subsection 4.1, the Θ space is more like the generation rules rather than a latent space, and we mention it for completeness.

3.2 GAN inversion methods

In the literature, there are mainly three kinds of GAN inversion methods, including optimization-based, learning-based, and hybrid methods. Table 3 [12, 14–16, 23, 35, 37, 38, 45, 46, 49–57, 59–65, 67, 90–93, 96–99, 101–160] also summarizes the characteristic of these methods, which shows the used backbone, latent space, inversion method, dataset, and their applications. In the following, the relevant methods are introduced in detail.

Optimization-based methods. Given a pre-trained GAN model, the objective of GAN inversion is to find an image \hat{x} generated by the GAN model which approximates a given image x, i.e.,

$$\hat{\boldsymbol{x}}^* = \arg\min_{\hat{\boldsymbol{x}} \in \mathcal{X}} \ell(\hat{\boldsymbol{x}}, \boldsymbol{x}), \tag{12}$$

where ℓ is a distance like ℓ_1 or ℓ_2 , and \mathcal{X} is the image space that is spanned by the images generated by the GAN model. We require the image to be in the GAN space in order to leverage the properties of GANs. Thus with (8), it can be rewritten as

$$\boldsymbol{z}^* = \arg\min_{\boldsymbol{z}\in\mathcal{Z}} \ell(G(\boldsymbol{z},\boldsymbol{\delta};\boldsymbol{\theta}_G),\boldsymbol{x}).$$
(13)

Here we use \mathcal{Z} space to illustrate the problem, and the formulation can be generalized to other latent spaces. Because the generator G is differentiable, Eq. (13) can be directly optimized via gradient descent algorithms [53, 91–93, 96–98, 104–107, 111–114, 118, 120, 124, 135], and the invertibility of GAN models is discussed by Aberdam et al. [120] and Ma et al. [107] with MLP-based and conv-based generators, respectively.

To promote the inversion accuracy and training stability, and boost the subsequent applications such as image editing and restoration, there are many improvements based on the basic formula in (13). Zhu et al. [103] adopted L-BFGS [161] algorithm instead of Adam [162]. Lipton et al. [105] and Shah et al. [106] proposed to use projected gradient descent (PGD) for better optimization in the latent space. Later, the inner loop of PGD was replaced by a neural network by Raj et al. [108]. Daras et al. [111] leveraged the attention map of the discriminator to boost the inversion. Huh et al. [118] and Kang et al. [98] further considered the geometric transformations of the subjects.

An obvious problem of optimization-based methods is efficiency. Since the optimization procedure usually requires thousands of update iterations, it takes a very long time for inversion. To eliminate this problem, Abdal et al. [91] proposed to initialize with the average latent code (i.e., \bar{w} obtained when training StyleGAN [14]). However, it still takes even minutes to invert a 1024 × 1024 image with a high-end GPU (e.g., NVIDIA Tesla V100). An alternative way is to leverage the learning-based methods, which predict the latent representation in a single forward pass.

Learning-based methods. Learning-based GAN inversion methods [52, 54-57, 90, 119, 123, 125, 128, 129, 131, 133, 134, 141] are generally equipped with an additional encoder E, and the latent code is optimized in an indirect way. Specifically, they optimize the parameters of the encoder, which is used to map the image back to the latent code, i.e., $E: x \to z$. Therefore, Eq. (13) can be re-written as

$$\boldsymbol{\theta}_{E}^{*} = \arg\min_{\boldsymbol{\theta}_{E}} \sum_{i} \ell(G(E(\boldsymbol{x}_{i};\boldsymbol{\theta}_{E}),\boldsymbol{\delta};\boldsymbol{\theta}_{G}),\boldsymbol{x}_{i}),$$
(14)

where $\boldsymbol{\theta}_E$ denotes the parameters of E.

It can be seen that the learning-based methods train the encoder with a whole dataset to learn the mapping from image to latent code, rather than optimize on a single image. Such a scheme yields several advantages. First, the optimization becomes smoother, which avoids the latent code from being trapped into a local optimum. Then, since the encoder has access to a large batch of images, it is easier for the encoder to learn some patterns from the data, e.g., the semantic directions for image manipulation in the latent space [52,90,119,123,129,131]. On the contrary, the optimization-based methods generally rely on

Table 3A summary of GAN inversion and methods leveraging pre-trained GANs for image editing and restoration. For the inversion method, "O", "L", "T" represent optimization-based, learning-based, and training-based (or fine-tuning) methods, while "-" means no inversion is performed in this method, and the numbers (without square brackets) are the indices of methods used for inversion in this table. Note that the methods are ordered (roughly) according to publicly accessible time (e.g., the appearance time on ArXiv, openreview.net, CVF Open Access, etc.)

No.	Method	Publication	Backbone	Latent space	Inversion method	Dataset*	$A_{\rm Pplication}^{\dagger}$
1	BiGAN [101]	ICLR 2017	-	Z	Т	MN, IN	Inv
2	ALI [102]	ICLR 2017	-	Z	т	CF, SV, CA, IN	Inv, Int
3	Zhu et al. [103]	ECCV 2016	DCGAN	Z	L, O	SH, LS, PL [‡]	Inv, Int, AE
4	IcGAN [90]	NeurIPSw 2016	cGAN	Z, C	L	MN, CA	Inv, AT, AE
5	Creswell et al. [104]	T-NNLS 2018	DCGAN, WGAN-GP	z	0	OM, UT, CA	Inv
07	DCD CAN [106]	ICLRW 2017	DCGAN	2	0	CA MN CA	Inv
8	Ma et al [107]	NeurIPS 2018	DCGAN	7	0	MN, CA	Inv IP
9	Suzuki et al. [49]	ArXiv 2018	SNGAN, BigGAN, StyleGAN	F	3	IN, FL, FF, DA	CO
10	GANDissection [45]	ICLR 2019	PGGAN	F	_	LS, AD	AE, AR
11	NPGD [108]	ICCV 2019	DCGAN, SAGAN	Z	L, O	MN, CA, LS	Inv, SR, IP
12	Image2StyleGAN [91]	ICCV 2019	StyleGAN	$\mathcal{W}+$	0	FF [‡]	Inv, Int, AE, ST
13	Bau et al. [109]	ICLRw 2019	${\rm PGGAN}, {\rm WGAN}\text{-}{\rm GP}, {\rm StyleGAN}$	$\mathcal{Z}, \ \mathcal{W}$	L, O	LS	Inv
14	GANPaint [50]	ToG 2019	PGGAN	\mathcal{Z}, Θ	$L,\;O,\;T$	LS	Inv, AE
15	InterFaceGAN [110]	CVPR 2020	PGGAN, StyleGAN	\mathcal{Z}, \mathcal{W}	3, 8	CH	AE, AR
16	GANSeeing [46]	ICCV 2019	PGGAN, WGAN-GP, StyleGAN	Z, W	13	LS	Inv
17	YLG [III]	CVPR 2020	SAGAN	2	0	IN LC FF	Inv
18	mage2StyleGAN++ [97]	CVPR 2020	StyleGAN PCCAN StyleCAN	7 VV+, N	0	LS, FF	INV, CO, IP, AE, SI
20	MimicGAN [112]	UCV 2020	DCGAN	7	0	CA FE LE	Inv. UDA AD AN
21	PULSE [53]	CVPB 2020	StyleGAN	z	0	FF. CH	Inv. SB
22	DGP [114]	ECCV 2020	BigGAN	z	О, Т	IN, P3	Inv, Int, IC, IP,
			0		,	,	SR, AD, TR, AE
23	StyleGAN2Distillation [115]	ECCV 2020	StyleGAN2, pix2pixHD	$\mathcal{W}+$	-	FF	AT, AE
24	EditingInStyle [116]	CVPR 2020	${\rm PGGAN},{\rm StyleGAN},{\rm StyleGAN2}$	\mathcal{F}	-	FF, LS	AT
25	StyleRig [51]	CVPR 2020	StyleGAN	$\mathcal{W}+$	-	FF	AT
26	ALAE [117]	CVPR 2020	StyleGAN	w	т	MN, FF, LS, CH	Inv, AT
27	IDInvert [38]	ECCV 2020	StyleGAN	$\mathcal{W}+$	L, O	FF, LS	Inv, Int, AE, CO
28	pix2latent [118]	ECCV 2020	BigGAN, StyleGAN2	Z	0	IN, CO, CF, LS	Inv, TR, AE
29	IDDistanglement [119]	ToG 2020	StyleGAN	W Z	L	FF	Inv, AT
30	Cuan at al [121]	ArXiv 2020	MLP StyleCAN	20	T O	MIN CH CD	Inv, IP
32	SeFa [122]	CVPB 2021	PGGAN BigGAN StyleGAN	Z	19 27	FF CH LS IN SS DA	AE
33	GH-Feat [123]	CVPR 2021	StyleGAN	s	L L	MN. FF. LS. IN	Inv. AT. AE
34	pSp [54]	CVPR 2021	StyleGAN2	\tilde{w} +	L	FF, AF, CH, CM	Inv, FF, SI, SR
35	StyleFlow [52]	ToG 2021	StyleGAN, StyleGAN2	$\mathcal{W}+$	12	FF, LS	AT, AE
36	PIE [124]	ToG 2020	StyleGAN	$\mathcal{W}+$	0	FF	AT, AE
37	Bartz et al. [125]	BMVC 2020	StyleGAN, StyleGAN2	$\mathcal{Z}, \ \mathcal{W} +$	L	FF, LS	Inv, DN
38	StyleIntervention [92]	ArXiv 2020	StyleGAN2	S	0	FF	Inv, AE
39	StyleSpace [93]	CVPR 2021	StyleGAN2	S	0	FF, LS	Inv, AE
40	Hijack-GAN [126]	CVPR 2021	PGGAN, StyleGAN	Z	-	CH	AE
41	NaviGAN [99]	CVPR 2021	pix2pixHD, BigGAN, StyleGAN2	Θ	[15]	FF, LS, CS, IN	AE
42	GLEAN [55]	CVPR 2021	StyleGAN	<i>vv</i> +	L	FF, LS	Inv, SR
43	CEPCAN [56]	AFAIV 2020	StyleGAN, StyleGAN2	P 142	T	FF, MF ⁺	Inv, IC, IP, SR
44	EnjoyEditing [127]	ICLB 2021	PGGAN StyleGAN2	7	12	FF CA CH P3 TB	Inv. AE
46	SAM [128]	ToG 2021	StyleGAN	\tilde{w}_{+}	L	CA. CH	AE
47	e4e [129]	ToG 2021	StyleGAN2	w_{\pm}	L	FF, CH, LS, SC	Inv, AE
48	StyleCLIP [130]	ICCV 2021	StyleGAN2	w+, s	47, O	FF, CH, LS, AF	AE
49	LatentComposition [131]	ICLR 2021	PGGAN, StyleGAN2	Z	L	FF, CH, LS	Inv, IP, AT
50	GANEnsembling [132]	CVPR 2021	StyleGAN2	$\mathcal{W}+$	L, O	CH, SC, PT	Inv, AT
51	ReStyle [133]	ICCV 2021	StyleGAN2	$\mathcal{W}+$	L	FF, CH, SC, LS, AF	Inv, AE
52	E2Style [134]	T-IP 2022	StyleGAN2	W+	L	FF, CH	Inv, SI, PI, AT, IP, SR, AE, IH
53	GPEN [57]	CVPR 2021	StyleGAN2	$\mathcal{W}+, \mathcal{N}$	L	FF, CH	Inv, SR
54	Consecutive [135]	ICCV 2021	StyleGAN	$\mathcal{W}+$	0	FF, RA	Inv, Int, AE
55	BDInvert [98]	ICCV 2021	StyleGAN, StyleGAN2	$\mathcal{F}/\mathcal{W}+$	0	FF, CH, LS	Inv, AE
56 57	HFGI [136]	CVPR 2022	StyleGAN2	vv+, F	L	FF, CH, SC	Inv, AE
57	Visualvocab [137]	CVPP 2021	BIGGAN StyleCAN2	ے ایم	т	P3, IN	AE Inv AE OT
59	GANGealing [139]	CVPR 2022	StyleGAN2	w	ц —	LS. FF. AF. CH. CU	TR.
60	HyperInverter [140]	CVPR 2022	StyleGAN2	<i>w</i> . Θ	L	FF, CH. LS	Inv, Int. AE
61	StyleGAN3-editing [141]	ArXiv 2022	StyleGAN3	s	L	FF, AF, LHQ	Inv, AE
62	InsetGAN [142]	CVPR 2022	StyleGAN2	w +	0	FF, DF [‡]	CO, IG
63	HairMapper [143]	CVPR 2022	StyleGAN2	$\mathcal{W}+$	47	FF, CM^{\ddagger}	AE
64	SAMInv [144]	CVPR 2022	BigGAN-deep, StyleGAN2	$\mathcal{W}+,\ \mathcal{F}$	L	FF, LS, IN	Inv, AE

64 SAMINV [144] CVPR 2022 BigGAN-deep, StyleGAN2 W+, F L FF, LS, IN Inv, AE
 ⁶⁴ Abbreviations: AD (ADE20K [145]), AF (AFHQ [23]), CA (CelebA [35]), CD (CACD [146]), CF (CIFAR [63]), CH (CelebA-HQ [12]), CM (CelebAMask-HQ [60]), CO (MS COCO [147]), CS (CityScapes [65]), CU (Caltech-UCSD Birds [148]), DA (Danbooru [149], aka Anime Faces), DF (DeepFashion [64]), FF (FFHQ [14]), FL (Flowers [150]), IN (ImageNet [59]), LF (LFW [151]), LHQ (Landscapes HQ [152]), LS (LSUN [37]), MF (MetFaces [16]), MN (MNIST [61]), OM (Omniglot [153]), P3 (Places365 [67]), PL (Places [154]), PT (Oxford-IIIT Pet [155], aka Cats and Dogs), RA (RAVDESS [156]), SC (Stanford Cars [157]), SS (Streetscape [158]), SV (SVHN [62]), TR (Transient [159]), UT (UT Zappos50K [160]).
 [†]Abbreviations: AD (Adversarial Defense), AE (Attribute Editing, i.e., w/o reference), AN (Anomaly Detection), AR (Artifacts Removal), AT (Attribute Transfer, i.e., w/ reference), CO (Image Crossover), [U]DA ([Unsupervised] Domain Adaptation), DN (Image Denoising), FF (Face Frontalization), IC (Image Colorization), IG (Image Guerration), IH (Information), ST (Style Transfer), TR (Transform and Random Jittering).
 [‡]Some custom datasets collected or regenerated by the authors are omitted since they are not publicly available or can be generated automatically based on current public datasets.

based on current public datasets

Liu M, et al. Sci China Inf Sci May 2023 Vol. 66 151101:11



Figure 3 (Color online) Illustration of GAN inversion methods. \boldsymbol{x} and $\hat{\boldsymbol{x}}$ are given real image and generated image, respectively. The red dotted line means supervision. It can be seen that the in-domain constraint [38] requires the generated image $\hat{\boldsymbol{x}}$ can be inverted back into the latent space. Here, \boldsymbol{z} is not restricted in \boldsymbol{Z} space, and may refer to more generic latent code (e.g., $\boldsymbol{w}, \boldsymbol{f}$, etc.). (a) Optimization-based method; (b) learning-based method; (c) hybrid method; (d) hybrid method (with in-domain constraint).

other methods (will be introduced in Subsection 4.2). Moreover, while the optimization-based methods can only be optimized with real images (the latent code of the synthetic images is already available), the learning-based methods can be trained with synthetic images. In this way, they directly leverage the ground-truth latent code [90], and Eq. (14) can be written as

$$\boldsymbol{\theta}_{E}^{*} = \arg\min_{\boldsymbol{\theta}_{E}} \sum_{i} \ell(E(G(\boldsymbol{z}_{i}, \boldsymbol{\delta}; \boldsymbol{\theta}_{G}); \boldsymbol{\theta}_{E}), \boldsymbol{z}_{i}).$$
(15)

Some methods also introduce extra models as auxiliary information or guidance. For example, Nitzan et al. [119] incorporated a face landmark detection model to control the pose of the generated face images. Many studies [54, 57, 119] introduced face recognition [163, 164] or contrastive learning [165] models as ID/content similarity metrics, while some of them [45, 134] also introduced semantic segmentation or face parsing models. Furthermore, vision-language models are also used for inversion and editing [130].

With the encoder, the model can get the latent code in a single forward pass, but the inversion is generally less accurate than optimization-based methods. Therefore, Wei et al. [134] proposed to use a multi-stage refinement scheme to keep both accuracy and efficiency, Alaluf et al. [133] introduced an iterative procedure which is initialized with the average face latent \bar{w} .

Hybrid methods. For both efficiency and accuracy, Zhu et al. [103] proposed the hybrid method that initializes the latent code for optimization via learning-based methods, which was followed by a series of studies [38, 50, 108, 121, 132]. Among these methods, we would like to highlight the method named IDInvert [38], which leverages the encoder for in-domain inversion. Specifically, an encoder is trained like the learning-based methods following (14), and the discriminator is also leveraged for training. Then, following the optimization-based methods, the latent code z_0 obtained by the encoder is used to initialize the optimization procedure, and then the learning objective can be formulated by

$$\boldsymbol{z}^* = \arg\min_{\boldsymbol{z}\in\mathcal{Z}} \ell(G(\boldsymbol{z},\boldsymbol{\delta};\boldsymbol{\theta}_G),\boldsymbol{x}) + \lambda \|\boldsymbol{z} - E(G(\boldsymbol{z},\boldsymbol{\delta};\boldsymbol{\theta}_G);\boldsymbol{\theta}_E)\|_2,$$
(16)

where λ is the balancing hyper-parameter. It can be seen from (16) that the second term requires that the image generated with the optimized latent code can be mapped back via the encoder. In other words, the optimization is guaranteed to be performed within the manifold supported by the encoder (as well as the generator), and it is called in-domain inversion. Please refer to Figure 3(d) for an illustration.

Guan et al. [121] leveraged the encoder during the optimization process in a collaborative way. In particular, the encoder (i.e., the embedding network in their study) and the iterator participate in each other's procedure. In each iteration, the encoder predicts a latent code as a reasonable initialization for the iterator, which obviously speeds up the optimization. Then, the iterator obtains a better latent code via optimization, which conversely serves as supervision for training the encoder. Both the encoder and the iterator achieve better performance/efficiency, and the method naturally supports online learning.

Apart from the above methods, some studies choose to learn an encoder together with the generator [101, 102, 117, 166], where the invertibility is guaranteed during the training procedure. Some other generative models are invertible by design [167, 168], however, the latent representation is typically a tensor of the same size as the image, which is sampled from a very simple distribution (e.g., Gaussian), making the latent representation less meaningful, and is beyond the scope of this study.

4 Analyzing and understanding large-scale GAN models

In this section, we first analyze GAN models from the perspective of image generation, and review the recent progress on understanding the neurons of pre-trained GANs. Then, we explore the current progress on latent space disentanglement, one of the key problems of image manipulation using pre-trained GANs.

4.1 Neuron understanding

4.1.1 Preliminary

Analogous to other neural networks, GAN [1] models are also trained in a data-driven manner, making them a black box which is hard to understand and interpret. To open the black box and better understand the image generation process in GAN generators, a series of studies [45, 46, 109, 169] proposed to disassemble the pre-trained GAN models and see how objects in the output images are generated.

In particular, as introduced in Subsection 2.2, the generation process in the generator G could be written as $\hat{x} = G(z)$. Following the main model paradigm in the deep learning era, a typical GAN model is composed of a stack of neural layers (e.g., an *N*-layer GAN model). Thus we can decompose the generator G into two parts at the *i*-th layer, and Eq. (8) can be rewritten as

$$\hat{\boldsymbol{x}} = G(\boldsymbol{z}) = (G_{i+1}^N \circ G_1^i)(\boldsymbol{z}) = G_{i+1}^N (G_1^i(\boldsymbol{z})),$$
(17)

where G_i is the *i*-th layer of G, and $G_a^b = G_b \circ G_{b-1} \circ \cdots \circ G_{a+1} \circ G_a$. Then the feature maps f_i at the *i*-th layer (i.e., $f_i = G_1^i(z)$) contains all information for generating the objects in the output images.

4.1.2 GAN dissection

For a certain convolution layer, the activation at each position is obtained by the sum-of-product operation between the previous features and the convolution kernels. Since each output channel has a particular kernel, we can see it as a neuron, which potentially controls the generation of some objects, and the discussions in this subsection are based on this observation.

Bau et al. [45] first built the relationship between the generated objects and the feature maps from f_i via network dissection techniques [170]. Specifically, for each object class c, a semantic segmentation model S_c is deployed to get the binary segmentation result $S_c(\hat{x})$. Then, the relationship between this class and a feature map $f_{i,u}$ (also known as a neuron) is evaluated by the spatial agreement between the thresholded feature map and the segmentation result, which is evaluated by the intersection-over-union (IoU) measure, i.e.,

$$\operatorname{IoU}_{u,c} = \frac{\mathbb{E}_{\boldsymbol{z}} |(\boldsymbol{f}_{i,u}^{\uparrow} > t_{u,c}) \wedge S_c(\hat{\boldsymbol{x}})|}{\mathbb{E}_{\boldsymbol{z}} |(\boldsymbol{f}_{i,u}^{\uparrow} > t_{u,c}) \vee S_c(\hat{\boldsymbol{x}})|}, \quad t_{u,c} = \operatorname{arg\,max}_t \frac{\boldsymbol{I}(\boldsymbol{f}_{i,u}^{\uparrow} > t; S_c(\hat{\boldsymbol{x}}))}{\boldsymbol{H}(\boldsymbol{f}_{i,u}^{\uparrow} > t; S_c(\hat{\boldsymbol{x}}))}, \quad (18)$$

 \wedge and \vee represent intersection and union operations, \uparrow denotes the upsampling operation, and $t_{u,c}$ is determined by maximizing the portion of the mutual information I in the joint entropy H.

In order to further verify the causal effects between the neurons and the generation results, after identifying the units which are relevant to a specific class of objects, Bau et al. [45] further ablated (or inserted) the units by setting the activation of the relevant neurons to zero (or a class-specific constant). Figure 4 shows the results of such modifications. It can be seen that the whole church appears after ablating the trees, and the grasses are added to the other image. The result indicates that (1) the generation is under the control of the relevant neurons, whose activation values determine the absence of the objects in the final generated image, and (2) the previously unseen objects may still be generated, but sheltered by other objects. Tousi et al. [169] further explored this phenomenon in multiple layers, and achieved a more accurate manipulation. This procedure could be regarded as an inverse procedure of the classification task to some extent, where the final prediction also shelters other classes, and we recommend [171] for an intuitive and interactive understanding.

4.1.3 Generation rule manipulation

When adding some objects in an image as stated in Subsection 4.1.2, it will fail to generate the desired objects somewhere (e.g., adding a door in the sky). Since the objects could be successfully inserted in other regions (e.g., adding a door on the wall), Bau et al. [172] asserted that there are some rules in the

Liu M, et al. Sci China Inf Sci May 2023 Vol. 66 151101:13



Figure 4 (Color online) Removing or adding particular classes via relevant unit discovery in GANDissection [45].



Figure 5 (Color online) Manipulating the generation rules via [172]. The two input images are modified via the same model.

generation process learned from the data, and they seek to manipulate the rules for verification and a general modification method at the model level.

Specifically, as shown in Figure 5 [172], suppose the rule in a generator is like $I_{\text{dome}} = G_{\text{church} \to \text{dome}}$ $(G_{z \to \text{church}}(z))$, where the generator could be regarded as an associative memory that stores a set of key-value pairs $\{(k_i, v_i)\}$, and is formulated by

$$\boldsymbol{v}_i \approx \boldsymbol{W} \boldsymbol{k}_i,$$
 (19)

we can modulate the parameter W (the associative memory) to reorganize the key-value pairs. For example, we could change the rule church \rightarrow dome to church \rightarrow tree, and the results are shown in Figure 5. The results clearly show the generation process, and help to open the black box of GAN models.

4.2 Semantic disentanglement and discovery

Understanding visual concepts is vital for many vision tasks like image editing and restoration, which helps the models to generate semantically consistent results. In this subsection, we review the relevant methods from two perspectives, i.e., semantic disentanglement and discovery. The former cares more about encouraging the independence of latent variables, i.e., a perturbation to any component of a latent variable should result in a change to the output that is distinct, while the latter focuses on discovering disentangled latent representations of high interpretability, i.e., a perturbation to one dimension should result in a change along exactly one meaningful attribute.

4.2.1 Semantic disentanglement

In early GAN models (e.g., DCGAN [3]), the concepts or attributes are generally entangled to a large extent in the rather simple latent space (e.g., standard Gaussian). Although some methods [90] achieve disentangled attributes by leveraging the conditional GANs [13,89] trained on datasets with class labels (e.g., ImageNet [59], LSUN [37]) or attribute annotations (e.g., CelebA [35], RaFD [173]), massive human efforts are required, which is unacceptable for all tasks. Thus a series of methods [174–180] are proposed for better disentanglement in GANs.

A route of methods achieves disentanglement by maximizing the mutual information. InfoGAN [175] is an early exploration on this task, which introduces an extra latent code c alongside the original random variable z. The model architecture is the same as (10), while the meaning of c is automatically learned via mutual information maximization when optimizing the GAN models, i.e.,

$$\min_{C} \max_{D} \mathcal{L}_{\text{InfoGAN}}(G, D) = \mathcal{L}_{\text{GAN}}(G, D) - \lambda I(G(\boldsymbol{z}, \boldsymbol{c}; \boldsymbol{\theta}_G); \boldsymbol{c}),$$
(20)

where I(x; y) denotes the mutual information between x and y, and \mathcal{L}_{GAN} is defined in (9). In particular, a regressor is trained to predict the input c. As such, the changes in c are required to influence the generation in a distinct way, which forces c to be traceable and disentangled. Zhu et al. [177] argued that the latent variation predictability can represent the disentanglement, and improved InfoGAN by predicting the relative changes between two generated images with only one dimension of the latent vector changed, i.e.,

$$\min_{G} \max_{D} \mathcal{L}_{\text{VP-GAN}}(G, D) = \mathcal{L}_{\text{GAN}}(G, D) - \lambda I(G(\boldsymbol{z}; \boldsymbol{\theta}_G), G(\boldsymbol{z} + \epsilon \boldsymbol{d}; \boldsymbol{\theta}_G); \boldsymbol{d}),$$
(21)

where d is a one-hot vector with the same dimension as z, and $\epsilon \in \mathcal{N}(0, 1)$. Note that the first metric of GAN disentanglement is proposed based on the observation. They further assumed that the interpretable representation should be spatially consistent, and improved the variation predictability in [178].

Another category focuses on the gradient of the GAN models with respect to the inputs. Ramesh et al. [174] focused on the coordinate directions in the latent space, and proposed a spectral regularizer to align the leading right-singular vectors of the Jacobian of GAN models w.r.t. the inputs with the coordinate axes, so that the trajectories are semantically meaningful. Peebles et al. [176] achieved disentanglement by encouraging the Hessian matrix of the GAN models w.r.t. the inputs to be diagonal, and the Hutchinson's estimator is leveraged to approximate the regularization. It is worth noting that the authors also showed the ability of the Hessian Penalty to "deactivate" redundant components. Wei et al. [179] required that when perturbing a single dimension of the network input, the change in the output should be independent (and uncorrelated) with those caused by the other input dimensions. Based on this assumption, they changed the assumption in [176] from $\frac{\partial}{\partial z_i} (\frac{\partial G}{\partial z_i}) = 0$ to $[\frac{\partial G}{\partial z_i}]^T \frac{\partial G}{\partial z_i} = 0$. He et al. [180] embedded a linear subspace in each layer of the GAN model by directly adding a group of basis U_i , where the *i*-th noise z_i is rewritten as $U_i L_i z_i + \mu_i$, and L_i is a diagonal matrix indicating the importance of basis vectors, μ_i denotes the origin of the subspace. They required that the basis is orthogonal, i.e., minimizing $||U_i^T U_i - I||_F^2$.

As introduced in Subsection 2.2, recent large-scale GAN models (especially the StyleGAN [14] series) achieve disentangled latent representations by sending the conditional information to the generator layers (see Figure 2), and the StyleGAN model maps the latent code into an intermediate latent space, i.e., W space via an MLP. A series of regularization terms are deployed for more stable training and better disentanglement, e.g., the mixing regularization and perceptual path length regularization. Please refer to Table 2 for more information. However, since the disentanglement is implicitly achieved in the W space, researchers also focused on finding meaningful semantic directions in the GAN models, which are introduced as follows.

4.2.2 Semantic discovery

To find meaningful semantic directions in the latent spaces, the supervised, semi-supervised, and unsupervised methods are introduced respectively as follows.

Supervised methods. The most intuitive way to leverage the manually annotated supervision is to model the semantic space explicitly via conditional GANs, where the attribute vector is delivered into the GAN model. Perarnau et al. [90] followed this intuition in IcGAN and adopted two encoders to map images into both the random noise space and the attribute space. However, this scheme is unavailable for unconditional GANs, which only synthesize images from random noise. Besides, the methods require massive human efforts for data annotation, which makes it inflexible and expensive.

Therefore, researchers sought to find indirect ways to leverage the annotations. For example, Shen et al. [110] assumed that there is a hyperplane in the latent space for binary attributes (e.g., male and female), which serves as the separation boundary between the opposite attributes. Thus they proposed to leverage the normal vector of the hyper-plane as the editing direction, and the support vector machine (SVM) is deployed for obtaining the hyper-plane. Refs. [115, 126, 127] followed a similar assumption. Wang et al. [126] trained a proxy model which can map the input noise to the attribute space, and the gradient of the proxy model w.r.t. the input noise is used as the non-linear editing direction. Viazovetskyi et al. [115] leveraged the pre-trained image attribute classifier to find the class centers of attributes in the latent space, and used the direction between different centers to represent the editing direction. Zhuang et al. [127] introduced a group of learnable editing directions d, and the pre-trained image attribute classifier is used to align each of the directions with an attribute.

Instead of learning editing directions, Abdal et al. [52] introduced a conditional normalizing flow model in \mathcal{W} -space of StyleGAN, which takes the attributes as conditions to map the latent \boldsymbol{w} back into the noise space. To edit the attributes, one can obtain the edited latent \boldsymbol{w}' by inverting the noise with new attributes. Besides, several methods focused on editing specific attributes in similar ways, for example, Alaluf et al. [128] sought to edit the age via an age regressor, and Goetschalckx et al. [47] tried to make the images more memorable with a memorability predictor.

Semi-supervised methods. The methods mentioned above still find the attribute directions directly related to the supervision (i.e., the manually annotated attributes), and some researchers proposed to edit some attributes with indirect supervision. As introduced in Subsection 4.1, Bau et al. [45] revealed the relationship between the GAN model neurons and the appearance of objects, which inspires a series of disentanglement methods to achieve spatially precise control [92,93,123]. Nitzan et al. [119] leveraged a pre-trained identification model and a landmark detection model, with which the attributes that can be described via landmarks (e.g., expression) are decomposed with the identity; thus the attributes can be transferred to other faces. Tewari et al. [51] shared a similar idea but employed a pre-trained threedimensional morphable face model (3DMM) to learn synthetic face image editing, which is later extended to real faces in [124]. Huh et al. [118] learned transformation of input shift, rescaling, etc., and estimated spatial transformation together with the latent variable. In particular, the spatial positions of the objects are obtained via a pre-trained object detection model. Xu et al. [135] bypassed the requirement of the pre-trained classifier by leveraging video data and optical flow. Jahanian et al. [181] applied data augmentation methods to the original data and trained the model in a self-supervised manner, but they are limited to these simple semantic features. Zhu et al. [182] achieved local editing by providing a bounding box, where the optimization is forced to only edit the given regions. It is worth noting that the large-scale vision-language models are introduced by Patashnik et al. [130], which leverages the common latent space for visual-linguist features, and combines the image-domain semantic discovery with the intrinsic semantic information in languages.

Unsupervised methods. Compared to supervised and semi-supervised methods, the unsupervised scheme has two main advantages. First, it avoids the need for manual labeling, and thus is more applicable to generalize to other categories, which greatly promotes the practical value of the methods. Second, the lack of supervision information also implies the lifting of restrictions, i.e., the models are encouraged to find new semantic features, which are not necessarily labeled by human annotators.

Some methods for semantic disentanglement introduced in Subsection 4.2.1 also support semantic discovery, which mainly considers the gradient w.r.t. inputs or mutual information. For example, Ramesh et al. [174] found trajectories corresponding to the principal eigenvectors, which is extended by Wang et al. [183] via further introducing the Riemannian geometry metric. Peebles et al. [176] leveraged their Hessian-based regularization term to identify interpretable directions in BigGAN's latent space. Voynov et al. [36] adopted similar strategies as [177], which learns a group of editing basis. By adding a perturbation to the latent code, the model is required to predict the perturbation given the original and modified image pairs. Tzelepis et al. [184] further argued that existing methods assume the latent directions are linear, which limits the disentanglement. Thus they introduced an RBF kernel to map the latent representations into a non-linear space, and learned a non-linear RBF path.

As analyzed in Subsection 4.1.2, the final output is determined by the features in the previous layers; thus GAN features can be considered for easier semantic discovery than the image domain. Following this idea, Härkönen et al. [48] proposed that the principal components of early layer features represent important factors of variations, so that they find semantic directions by PCA operation over the early feature space. Collins et al. [116] sought to find attributes relevant to a region of interest (ROI), and proposed to cluster the features in a layer and sort the channels by the relevance with the ROI regions (following the idea of [45] but in an unsupervised way). From another perspective, the parameters also play an important role in the generation process, which has also been introduced in Subsection 4.1.3. Therefore, there are also some methods leveraging the parameter space for semantic discovery. Shen and Zhou [122] formulated the perturbation in a layer and the corresponding changes in the output image, and finally the relationship between the semantic directions and the parameter matrix is derived, leading to a closed-form solution, which achieves an efficient and effective semantic factorization method (termed by SeFa). Cherepkov et al. [99] followed similar strategies to [36, 183], and proposed two methods in the parameters.



Figure 6 (Color online) Out-of-distribution image interpolation. Before interpolation, two cat images from the AFHQ dataset [23] are mapped back to the latent space (i.e., z_1 and z_2) via a GAN model [91] trained with the FFHQ dataset [14]. It can be seen that even the image can be well reconstructed, the interpolation results ($G(\lambda \cdot z_1 + (1 - \lambda) \cdot z_2)$) are looming the shape of human faces.

5 Applications to image editing and restoration

Since most GAN models are trained in an unsupervised manner, they are not limited to a specific task and can serve as a generic prior for many other tasks with visual outputs (e.g., images). In this section, we will go through these applications, where a brief summary can be found in Table 3.

5.1 Image editing

As introduced in Section 3 and Subsection 4.2, the latent spaces of recent GAN models have been widely explored in the literature. With the exhilarating attribute disentanglement ability of the latest GAN models, we are able to discover and traverse the latent spaces in a more semantic-aware way, making image editing applicable with pre-trained GAN models. In this subsection, we briefly review the image editing applications relying on semantic disentanglement and discovery methods, including image interpolation, style transfer, image crossover, attribute editing, image transform, etc.

5.1.1 Image interpolation (image morphing)

Image interpolation, also known as image morphing, aims at interpolating two images semantically. By observing the generated images, we can have an intuitive sense of the space used for interpolation; i.e., a well-characterized space should be reflected in smooth and semantically appropriate transitions in the interpolated images. Therefore, apart from generating new images, image interpolation is often utilized to evaluate the quality of the discovered latent space. For example, as shown in Figure 1, the interpolation operation in the pixel space (i.e., $\mathbf{I} = \alpha \cdot \mathbf{x}_1 + (1 - \alpha) \cdot \mathbf{x}_2$) leads to obvious artifacts, and the results cannot meet the semantic expectations. On the contrary, the interpolation in the latent space (e.g., [38, 91, 103, 135]) is able to well describe the change of view, which can be formulated by

$$\boldsymbol{I} = \boldsymbol{G}(\boldsymbol{\beta} \cdot \boldsymbol{z}_1 + (1 - \boldsymbol{\beta}) \cdot \boldsymbol{z}_2); \boldsymbol{\theta}_G),$$
(22)

where z_1 and z_2 can be obtained by the inversion methods (Section 3) from two images x_1 and x_2 .

There are two special cases that are worth mentioning. As revealed by Abdal et al. [91], when performing GAN inversion via optimization-based methods, a GAN model trained with face datasets can well reconstruct an image from other classes (e.g., cats). However, since the optimization-based methods are less constrained and can have a much larger accessible latent space, interpolating between these reconstructed images will not produce meaningful results (see Figure 6), which further verifies the propositions of IDInvert [38] about domain-regularized optimization (see (16)). Another thing is about the training/fine-tuning based inversion methods (see Table 3). For more accurate reconstruction, some of these methods (e.g., [114]) will derive a specific group of model parameters for each sample. In this way, there will be two sets of θ_G when performing image interpolation operations, which is inconsistent with (22). Fortunately, Wang et al. [185] showed that when fine-tuned from the same checkpoint, the parameters of two models can be interpolated to achieve an intermediate function. Thus, the image interpolation can be performed as

$$\boldsymbol{I} = G(\beta \cdot \boldsymbol{z}_1 + (1-\beta) \cdot \boldsymbol{z}_2); \beta \cdot \boldsymbol{\theta}_{G_1} + (1-\beta) \cdot \boldsymbol{\theta}_{G_2}),$$
(23)

where θ_{G_1} and θ_{G_2} are the specialized parameters for reconstructing x_1 and x_2 , respectively.



Figure 7 (Color online) Style transfer and image crossover. The images are processed by IDInvert [38].

5.1.2 Style and attribute transfer, image crossover

As shown in [14], the features in earlier layers control coarser-grained attributes like structure and pose, while the ones in latter layers control finer-grained attributes like appearance and textures. For the StyleGAN series, the latent codes are corresponding to these layers in the W space, and with the style mixing regularization during training, the latent codes in different layers can be distinct [14]. Therefore, the pre-trained StyleGAN models are ready-to-use for style transfer tasks by replacing the latent codes in latter layers with the ones from the style image [91,115,121,123,134]. In this way, the characters regarded as style (e.g., color, stroke, texture) are transferred to the content image. Instead of simply mixing style codes from different images, some studies also perform a more precise control by specifying the attributes to be transferred [51,119], which can be combined with the methods introduced in Subsection 4.2 for disentangling the attributes. Some methods [54, 134] also considered image-to-image translation tasks such as transferring sketch or parsing (segmentation) images into natural ones, or generating toonified versions of a human face.

Apart from the global tasks of style and attribute transfer, some methods focus on local modifications from two perspectives, i.e., local style transfer and image crossover. The former is the style transfer task in a local region, while the latter is to combine multiple images and regenerate the stitched image to make it visually plausible. For example, Ref. [131] introduced random masks (e.g., m_1 and m_2) during training, such that the encoder will be compatible with masked input images, making it possible to combine multiple incomplete images by corresponding mask directly,

$$\boldsymbol{I} = G(E(\boldsymbol{m}_1 \circ \boldsymbol{x}_1 + \boldsymbol{m}_2 \circ \boldsymbol{x}_2)), \tag{24}$$

which generates more realistic and natural images than blending in pixel or latent spaces like,

$$I = G(E(\alpha \cdot x_1 + (1 - \alpha) \cdot x_2)), \qquad (pixel space blend) \qquad (25)$$

$$I = G(\alpha \cdot E(x_1) + (1 - \alpha) \cdot E(x_2)).$$
 (latent space blend) (26)

Note that Image2StyleGAN++ [97] performs local style transfer in both W+ and N spaces. Ans some methods transfer the features rather than latent codes [49,116]. Visual results are given in Figure 7.

5.1.3 Attribute editing

Attribute editing aims to modify certain attributes of an image (e.g., age, gender), which is a popular application of GAN models in recent years. The difference between attribute editing and attribute transfer in Subsection 5.1.2 is that the reference image is not necessary for attribute editing tasks. Some methods

Liu M, et al. Sci China Inf Sci May 2023 Vol. 66 151101:18



Figure 8 (Color online) Attribute editing with manipulation directions found by InterFaceGAN [110].

focus on physical attributes. For example, Zhu et al. [103] and Abdal et al. [97] took color strokes provided by users as input and optimized the image towards a reasonable appearance where the color or shape of regions covered by the strokes is modified, while the background regions are kept unchanged. Jahanian et al. [181] focused on self-supervised training for zooming, shifting, rotating, brightness adjustment, and etc.

To manipulate the attributes more closely related to semantic concepts, some methods [47, 52, 90, 114, 115, 128] leveraged pre-trained attribute classifiers or attribute labels of the datasets to identify the manipulation direction in the latent space. Some methods [38,93,110,112] chose to first define an attribute separation hyperplane and take the normal vector as editing directions. Considering the non-linearity of the latent space, especially when manipulating the attributes to a large extent, Wang et al. [126] proposed to learn a proxy model to estimate the gradient of the pre-trained GAN models and modify the attributes step by step. The primary drawback of the aforementioned methods is that they need annotated labels (or classification models trained with these labels) for training. To eliminate the dependence on these labels, some studies [99,122] proposed unsupervised methods to disentangle different attributes, which have been introduced in detail in Subsection 4.2. Besides, some methods [45,50,92] achieved local attribute editing based on the observations introduced in Subsection 4.1.2.

There are also some other interesting methods to achieve attribute editing. Image2StyleGAN [91] used a one-shot method that obtains the manipulation direction vector via the difference between two images. Ref. [123] sampled global or local features in the latent space which resulted in random new attributes. E4E [129] leveraged an encoder for generating target latent codes directly instead of the previous inversion+editing methods. Alaluf et al. [141] leveraged StyleGAN3 for image editing, and designed an encoder network for images that are unaligned. Visual results of image editing are given in Figure 8.

5.1.4 Image transform and jittering, visual alignment

Since StyleGAN [14] introduces layer-wise noises (i.e., the \mathcal{N} space), it naturally supports random jittering on different layers. Pan et al. [114] further added Gaussian noise to the \mathcal{Z} space of BigGAN model [13] and achieved image jittering with larger variance. Jahanian et al. [181] also explored the image transformation with pre-trained GANs. As introduced in Subsection 2.2, StyleGAN3 [17] adopts equivariant layers, which greatly facilitates the output image quality when performing transform operations like shifting.

Based on the image transform operations, some methods leverage pre-trained GANs for visual alignment. For example, pSp [54] constrained that face images are inverted to the same latent code with its horizontally flipped counterpart, which leads to a frontal face for profile ones. Peebles et al. [139] further extended the method to more general cases by clustering one or several center points (e.g., frontal for faces), and finally they can obtain a spatial transformer network that can perform visual alignment as well as the reverse process. In other words, we can edit the transformed (aligned) images for dealing with a certain object in all video frames. In Figure 9, we show several examples of these applications.

5.2 Image restoration

Unlike the applications introduced in Subsection 5.1, which mainly leverage the exhilarating latent space disentanglement ability of recent GAN models, image restoration tasks rely more on the high-quality natural image priors provided by the pre-trained large-scale GAN models. Note that the GAN model



Figure 9 (Color online) Dense visual alignment method [139] trained with GAN supervision. Please zoom in for better observation.

is trained to learn the mapping $g : z \mapsto x$, where z is generally sampled from a simple distribution (e.g., the standard Gaussian distribution $z \sim \mathcal{N}(0, I)$), while $x \in \mathcal{X}$ represents the high-quality training datasets. With the millions or billions of training iterations, we assume that $\{z_i\}$ covers at least a truncated Gaussian space, or more precisely, a manifold, which is defined by \mathcal{M} . In other words, given a latent code sampled in \mathcal{M} , a well-trained GAN model helps to constrain the output to the \mathcal{X} space; i.e., the output tends to be a high-quality image following the training image distribution. In this way, given a corrupted image, we can find a latent code in the latent space, which has the smallest "distance" from this corrupted image. Then the image generated from the latent code can be regarded as the one with the highest likelihood among all possible high-quality counterparts. The relevant image restoration applications include image super-resolution, image denoising, image inpainting, image colorization, and artifacts removal, and we will go through these applications in this subsection.

5.2.1 Image super-resolution, image denoising, image inpainting, and image colorization

Considering the common methods used for image super-resolution, image denoising, image inpainting, and image colorization, we introduce these applications together in this part. Overall, these methods can be divided into two classes, i.e., unsupervised and supervised. Specifically, for unsupervised methods [53, 112, 114], the optimization objective is typically defined by

$$\mathcal{L}_{\text{unsup}} = \arg\min_{\boldsymbol{z}} \|\zeta(G(\boldsymbol{z};\boldsymbol{\theta}_G)) - \boldsymbol{x}_{\text{LQ}}\|_p^p + \rho(\boldsymbol{z}),$$
(27)

where ζ denotes the degradation function (e.g., down-sampling for image super-resolution, color-to-gray for image colorization, masking for image inpainting), and ρ denotes the regularization on z derived from the distribution of the latent space (e.g., Gaussian). Here we note that, as shown in Figure 10, the denoising task here is actually inpainting tasks with random defective pixels rather than additive noise suppressing, and for super-resolution tasks the degradation process is approximated by simple downsampling algorithms like bicubic. These drawbacks make the optimization-based unsupervised methods limited to a certain range of tasks, where the ground-truth images are unavailable during inference.

To solve this problem, with the development of learning-based methods, the ground-truth (reference) images can be utilized for training the encoder to obtain the restoration results, i.e., the supervised methods. With the reference, the encoder can directly learn the mapping $f : \mathbf{x}_{LQ} \mapsto \mathbf{z}_{HQ}$, which is then delivered into the pre-trained GAN model to obtain the high-quality output image $\mathbf{x}_{HQ} = G(\mathbf{z}_{HQ}; \boldsymbol{\theta}_G)$ [54–57,96,125,134]. In particular, the loss function is defined by

$$\mathcal{L}_{sup} = \sum_{i} \|\xi_{i}(G(E(\boldsymbol{x}_{LQ};\boldsymbol{\theta}_{E});\boldsymbol{\theta}_{G}) - \xi_{i}(\boldsymbol{x}_{HQ})\|_{p}^{p} + \rho(\boldsymbol{z}),$$
(28)

where ξ_i denotes the identity operation or models for loss functions like identity loss [163], LPIPS loss [82], and face parsing loss [60], which are utilized for better identity preserving or perceptual image quality. It is worth noting that, unlike other methods which produce the output images directly by the pre-trained



Figure 10 (Color online) Image restoration with GPEN [57] pre-trained with FFHQ [14] dataset.



Figure 11 (Color online) Artifacts removal via GANDissection [45], where the neurons controlling the generation of artifacts are set to zero.

GAN model, GLEAN [55] stacks an extra decoder on the top of the GAN model, where the features from the pre-trained GAN model as taken as intermediate features to improve the quality of the output generated by the extra decoder. Visual results can be found in Figure 10.

5.2.2 Artifact removal

As shown by Bau et al. [45], similar to other objects, the generation of some artifacts is also controlled by the neurons of the GAN model. Therefore, the network dissection technique can also be used to identify the neurons that cause the artifacts, which can bring an obvious image quality improvement (the FID [75] improves from ~ 43 to ~ 27). Shen et al. [110] further explored artifact removal in the latent space, where a linear-SVM is trained with 4k bad results to obtain the separation hyperplane, and the artifacts are successfully suppressed by using the normal vector. Ref. [169] trained an artifact classifier with a proposed dataset, and identified the flawed regions via Grad-CAM [186]. Then the artifacts are suppressed in a sequential process. The visual results are shown in Figure 11.

5.3 Other applications

Apart from the aforementioned image editing and restoration applications, the pre-trained GAN models can also be applied to other tasks, for example, information hiding [134], unsupervised domain adaptation [113], adversarial defense [113,114], anomaly detection [113], 3D reconstruction [187,188], and so on. Besides, a new trend in leveraging pre-trained GAN models is to combine multiple GAN models for generating different parts (e.g., the face region and the whole body) [142], which has great potential for generating images with complex scenes and rich contents.

6 Discussion

6.1 Comparison with other generative models

Apart from GANs, there are several well-known generative models, e.g., variational auto-encoder (VAE) [189], normalizing flow [167], autoregressive models [190], and diffusion models [168]. In this subsection, we briefly compare GANs with these models from several aspects.

Generation quality. In terms of generation quality, GAN models and diffusion models perform best. Recently, diffusion models start to emerge and have obtained a lot of attention. They are also successfully applied in several large-scale text-guided image generation models like DALL·E 2 [191] and Imagen [192].

Inference speed. However, diffusion models have a major problem with inference speed, where tens or hundreds of steps are required. By contrast, VAE, normalizing flow, and GANs can generate images in a single forward pass. Autoregressive models are the least efficient due to pixel-by-pixel generation.

Diversity in generated images. In terms of diversity, the mode collapse problem of GAN models is well known. The other generation models have better performance than GANs on diverse generations.

Semantic representations. Thanks to the structural design of the StyleGAN series, GAN models have the best semantic representation in the latent space, which has been shown in Subsection 4.2. Normalizing flow and diffusion models require that the latent code should have the same size as the image, which limits their semantic representation ability in the latent space.

To sum up, GAN models have been extensively explored, and the well-constructed latent representation makes it suitable for semantic-oriented tasks. The diffusion models have shown appealing image generation ability, yet utilizing pre-trained diffusion models for subsequent tasks remains to be explored.

6.2 Connections to high-level vision tasks

In this survey, we discuss the applications of pre-trained GAN models on low-level vision tasks like image editing and restoration, which largely leverage the generative image priors learned by the GAN models. Since GANs have also been deployed in high-level vision tasks [193–195], for completeness, we also briefly discuss the connections between pre-trained GANs and high-level vision tasks. It is noting that high-level tasks are more closely related to the image generation process introduced in Subsection 4.1.2, which makes the features of pre-trained GAN models more discriminative for high-level vision tasks (e.g., part segmentation).

Pre-trained GANs for data augmentation. The most intuitive way to utilize pre-trained GAN models for high-level vision tasks is generating extra samples for data augmentation [196, 197]. The GAN-based data augmentation can help with expanding the data scale, generating out-of-distribution samples, and balancing the data distribution between classes.

High-level tasks with pre-trained GANs. As shown in [198], the features of pre-trained GAN models are "discriminative" enough for part segmentation tasks. Ref. [199] further proposed a progressive finetuning strategy to reduce the gap between image generation and part segmentation. For anomaly detection, both the generator and the discriminator were employed for training [200]. Besides, Refs. [201, 202] leveraged pre-trained GAN models for initializing adversarial attack models.

6.3 Challenges and open problems

6.3.1 GAN inversion

For leveraging pre-trained GAN models on real images, it is a fundamental task to invert them into the latent space, with which the images should be well reconstructed via the GAN models. Although existing GAN inversion methods have demonstrated decent performance, they are still far from perfect, and there are many challenging problems that remain unsolved.

Fidelity and editability. When inverting a real image into latent space, there exists a trade-off between the reconstruction fidelity and latent code editability. On the one hand, for methods that only use the \mathcal{W} space (or the derived spaces like $\mathcal{W}+$ or \mathcal{S} spaces), their reconstruction accuracy is greatly limited due to the narrow latent space in comparison to the image space. Besides, mode dropping [46] is a severe problem influencing the inversion quality. On the other hand, for methods that leverage spatial dimensions (e.g., the \mathcal{F} [98], Θ [99,138], and \mathcal{N} [57,97] spaces), the editability becomes a main concern as these spaces are less disentangled. Wang et al. [136] took a step forward by controlling the passage of spatial-dimension features with the latent code in \mathcal{W} space, which achieves a better trade-off between fidelity and editability. However, it still remains to be better solved for real-world applications, especially when there are geometry changes (e.g., zoom, rotation) or camera movements.

Out-of-distribution generalization. Since mode dropping [46] occurs even when a subject has shown in the training images, the out-of-distribution generalization ability (on unseen classes or domains) is a main concern of pre-trained GANs. Indeed, with less constrained optimization, the GAN models are able to generate images from classes that have no intersection with the training set [91]. Nevertheless, the resultant latent code does not fall into the meaningful latent space; i.e., the semantic directions in the latent space are inconsistent with the out-of-distribution images. So that we are unable to edit these images following the prior learned in the training set. Some methods have tried to adapt GAN models to new domains via few-shot learning via cross-domain correspondence [203], but the properties like editability require further exploration.

6.3.2 Image generation, editing, and restoration

Current methods have achieved appealing performance on image generation, editing, and restoration tasks, where the results are photo-realistic and the semantic information is well represented. Here we conclude some potential directions for better leveraging pre-trained GANs.

Multi-GAN composition. Due to the limited model capacity and training mechanism, it is rather hard for a single GAN model to generate a complex scene with every object photo-realistic and natural. Therefore, Frühstück et al. [142] tried to combine two pre-trained GAN models together for a compositional GAN model, one of which is for generating the whole body of a person, while the other is responsible for providing a better face region. The two GAN models are combined by joint-optimizing the latent codes. It is of great potential to generate complex scenes via the combination of multiple GAN models.

Physical rules and 3D awareness. For the purpose of image editing or video generation, the GAN models should be aware of the physical rules for realistic results, and the 3D awareness also benefits to produce reasonable structures. As introduced in Subsection 5.1.4, the pre-trained GAN models are able to manipulate the geography properties of the images. In other words, they have implicitly learned some physical rules from the training set. Besides, the 3D priors like 3DMM have been implied for learning face editing models [51, 124], and neural radiance field (NeRF) has also been incorporated in some 3D-aware generative models [204]. These models have shown the great advantage of 3D awareness. It is also worth noting that the embedded physical rules and 3D awareness will benefit the combination of multiple GAN models for generating complex scenes.

Diverse and reliable results. A major problem of existing image editing and restoration methods based on pre-trained GANs is that only one result is obtained for a given input. Yet these two tasks are both ill-posed, which prefer diverse results for more flexible choices. More importantly, for image restoration tasks, since the input image is usually greatly degraded, the results are sometimes less reliable. In this condition, we argue that providing multiple results with confidence scores and/or precise editing directions recommended based on the degraded input will be more applicable, where timely feedback and adjustment can be obtained through user participation (e.g., eyewitness).

Multi-modality combination. It is common sense that the generated images should be semantically consistent, and we hope that we can directly provide the semantic descriptions to modify the generated (edited, restored) results. To this end, many studies [130, 191, 192] have sought to combine language models with generative models, and exhibit extraordinary performance in text-guided image editing. It is challenging yet exciting for the pre-trained GAN models to be well-aligned with semantic concepts in the languages, which will make the real-world application based on pre-trained GANs more convenient.

Toward more general image restoration. Even with well-trained GAN models as prior, image restoration is still a challenging ill-posed problem, and the GAN prior is restricted by the training data. To alleviate this problem, some studies [205] discussed the possibility of face image restoration without face and GAN priors, which are limited to simple degradations (e.g., bicubic interpolation) but are worth further exploration. Li et al. [206] proposed to solve this problem by considering the trade-off of accuracy and universality between GAN models trained with faces images or natural images, and explored to transfer priors of faces to natural images by learning degradation model from the face regions. AirNet [207] proposed to extract degradation representations via contrastive learning from the input images, which can also be combined with pre-trained GAN models for better blind restoration. In a word, it still requires tremendous efforts toward more general image restoration.

6.3.3 Others

Combining with other generative models. Compared to other generative methods (e.g., VAE [189], Normalizing Flow [167], and Diffusion Models [168]), the advantages of GAN models lie in the superior image quality (vs. VAEs), flexible structure (vs. normalizing flows), and efficient sampling (vs. diffusion models). However, GAN models approximate the data distribution in an implicit way, and they are unstable to train and can easily lead to mode collapse. A potential way is to combine the GAN models with other generative methods, which can give full play to the advantages of both methods. For example, Lyu et al. [208] proposed to use GANs to generate an intermediate result for diffusion models, Grover et al. [209] combined the maximum likelihood of normalizing flow with GANs.

Evaluation metrics. For boosting the development of GANs and techniques of leveraging pretrained GAN models, the evaluation metrics play an important role. For intuitive evaluations such as reconstruction quality and mode dropping, there have been some methods like FID [75] and FSD [46]. However, indirect or more complex evaluations like the detail and semantic consistency after editing and restoration are less explored. Besides, for the aforementioned challenging problems, corresponding evaluation metrics are also desired to evaluate the effectiveness.

6.4 Conclusion

Behind the superior generation ability, GAN models have shown the concept abstraction ability and the learned generative image priors, which have been extensively explored in recent methods. In this study, we give a comprehensive survey of recent progress on leveraging pre-trained GAN models for image editing and restoration tasks, where the two main features of recent GAN models are utilized, i.e., the disentanglement ability of the latent space and the generative image priors inherently learned by the GAN models. By introducing the neuron understanding methods, we try to construct an intuitive and in-depth impression of pre-trained GAN models. Subsequently, we review the GAN inversion methods as well as the semantic disentanglement and discovery techniques, and the relevant applications of image editing and restoration tasks are introduced. Finally, we discuss some challenges and open problems in leveraging pre-trained GAN models for image editing and restoration.

Acknowledgements This work was supported by National Natural Science Foundation of China (Grant Nos. U19A2073, 62006064), Hong Kong RGC RIF (Grant No. R5001-18), and 2020 Heilongjiang Provincial Natural Science Foundation Joint Guidance Project (Grant No. LH2020C001).

References

- 1 Goodfellow I, Pouget-Abadie J, Mirza M, et al. Generative adversarial nets. In: Proceedings of International Conference on Neural Information Processing Systems, 2014
- 2 Denton E L, Chintala S, Fergus R, et al. Deep generative image models using a laplacian pyramid of adversarial networks. In: Proceedings of International Conference on Neural Information Processing Systems, 2015
- 3 Radford A, Metz L, Chintala S. Unsupervised representation learning with deep convolutional generative adversarial networks. In: Proceedings of International Conference on Learning Representations, 2016
- 4 Zhang H, Xu T, Li H, et al. StackGAN: text to photo-realistic image synthesis with stacked generative adversarial networks. In: Proceedings of IEEE International Conference on Computer Vision, 2017. 5907–5915
- 5 Zhang H, Goodfellow I, Metaxas D, et al. Self-attention generative adversarial networks. In: Proceedings of International Conference on Learning Representations, 2019. 7354–7363
- 6 Mao X, Li Q, Xie H, et al. Least squares generative adversarial networks. In: Proceedings of IEEE International Conference on Computer Vision, 2017. 2794–2802
- 7 Berthelot D, Schumm T, Metz L. BEGAN: boundary equilibrium generative adversarial networks. 2017. ArXiv:1703.10717
- 8 Jolicoeur-Martineau A. The relativistic discriminator: a key element missing from standard gan. In: Proceedings of International Conference on Learning Representations, 2019
- 9 Arjovsky M, Chintala S, Bottou L. Wasserstein generative adversarial networks. In: Proceedings of International Conference on Learning Representations, 2017. 214–223
- 10 Gulrajani I, Ahmed F, Arjovsky M, et al. Improved training of wasserstein GANs. In: Proceedings of International Conference on Neural Information Processing Systems, 2017. 5769–5779
- 11 Miyato T, Kataoka T, Koyama M, et al. Spectral normalization for generative adversarial networks. In: Proceedings of International Conference on Learning Representations, 2018
- 12 Karras T, Aila T, Laine S, et al. Progressive growing of GANs for improved quality, stability, and variation. In: Proceedings of International Conference on Learning Representations, 2018
- 13 Brock A, Donahue J, Simonyan K. Large scale GAN training for high fidelity natural image synthesis. In: Proceedings of International Conference on Learning Representations, 2018
- 14 Karras T, Laine S, Aila T. A style-based generator architecture for generative adversarial networks. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2019. 4401–4410
- 15 Karras T, Laine S, Aittala M, et al. Analyzing and improving the image quality of StyleGAN. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 8110–8119
- 16 Karras T, Aittala M, Hellsten J, et al. Training generative adversarial networks with limited data. 2020. ArXiv:2006.06676
 17 Karras T, Aittala M, Laine S, et al. Alias-free generative adversarial networks. In: Proceedings of International Conference on Neural Information Processing Systems, 2021
- 18 Isola P, Zhu J Y, Zhou T, et al. Image-to-image translation with conditional adversarial networks. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2017. 1125–1134
- 19 Zhu J Y, Park T, Isola P, et al. Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of IEEE International Conference on Computer Vision, 2017. 2223–2232
- 20 Choi Y, Choi M, Kim M, et al. StarGAN: unified generative adversarial networks for multi-domain image-to-image translation. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2018. 8789–8797
- 21 He Z, Zuo W, Kan M, et al. AttGAN: facial attribute editing by only changing what you want. IEEE Trans Image Process, 2019, 28: 5464–5478
- 22 Liu M, Ding Y, Xia M, et al. STGAN: a unified selective transfer network for arbitrary image attribute editing. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2019. 3673–3682
- 23 Choi Y, Uh Y, Yoo J, et al. StarGAN v2: diverse image synthesis for multiple domains. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 8188–8197
- Ledig C, Theis L, Huszár F, et al. Photo-realistic single image super-resolution using a generative adversarial network.
 In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2017. 4681–4690

Liu M, et al. Sci China Inf Sci May 2023 Vol. 66 151101:24

- 25 Wang X, Yu K, Wu S, et al. ESRGAN: enhanced super-resolution generative adversarial networks. In: Proceedings of European Conference on Computer Vision, 2018
- 26 Wang X, Xie L, Dong C, et al. Real-ESRGAN: training real-world blind super-resolution with pure synthetic data. In: Proceedings of IEEE International Conference on Computer Vision Workshops, 2021. 1905–1914
- Zhang K, Liang J, van Gool L, et al. Designing a practical degradation model for deep blind image super-resolution.
 In: Proceedings of IEEE International Conference on Computer Vision, 2021. 4791–4800
- 28 Kupyn O, Budzan V, Mykhailych M, et al. DeblurGAN: blind motion deblurring using conditional adversarial networks. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2018. 8183–8192
- 29 Kupyn O, Martyniuk T, Wu J, et al. DeblurGAN-v2: deblurring (orders-of-magnitude) faster and better. In: Proceedings of IEEE International Conference on Computer Vision, 2019. 8878–8887
- 30 Zheng S, Zhu Z, Zhang X, et al. Distribution-induced bidirectional generative adversarial network for graph representation learning. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 7224–7233
- 31 Zhu H, Peng X, Chandrasekhar V, et al. DehazeGAN: when image dehazing meets differential programming. In: Proceedings of International Joint Conference on Artificial Intelligence, 2018. 1234–1240
- 32 Zhu H, Cheng Y, Peng X, et al. Single-image dehazing via compositional adversarial network. IEEE Trans Cybern, 2019, 51: 829–838
- 33 Mehta A, Sinha H, Narang P, et al. HiDeGAN: a hyperspectral-guided image dehazing GAN. In: Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 2020. 212–213
- 34 Dong Y, Liu Y, Zhang H, et al. FD-GAN: generative adversarial networks with fusion-discriminator for single image dehazing. In: Proceedings of AAAI Conference on Artificial Intelligence, 2020. 10729–10736
- 35 Liu Z, Luo P, Wang X, et al. Deep learning face attributes in the wild. In: Proceedings of IEEE International Conference on Computer Vision, 2015. 3730–3738
- 36 Voynov A, Babenko A. Unsupervised discovery of interpretable directions in the GAN latent space. In: Proceedings of International Conference on Learning Representations, 2020. 9786–9796
- 37 Yu F, Seff A, Zhang Y, et al. LSUN: construction of a large-scale image dataset using deep learning with humans in the loop. 2015. ArXiv:1506.03365
- 38 Zhu J, Shen Y, Zhao D, et al. In-domain GAN inversion for real image editing. In: Proceedings of European Conference on Computer Vision, 2020. 592–608
- 39 Rudin L I, Osher S, Fatemi E. Nonlinear total variation based noise removal algorithms. Physica D-Nonlinear Phenomena, 1992, 60: 259–268
- 40 Buades A, Coll B, Morel J M. A non-local algorithm for image denoising. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2005. 60–65
- 41 Elad M, Aharon M. Image denoising via sparse and redundant representations over learned dictionaries. IEEE Trans Image Process, 2006, 15: 3736–3745
- 42 Li B, Gou Y, Gu S, et al. You only look yourself: unsupervised and untrained single image dehazing neural network. Int J Comput Vis, 2021, 129: 1754–1767
- 43 Shoshan A, Mechrez R, Zelnik-Manor L. Dynamic-Net: tuning the objective without re-training for synthesis tasks. In: Proceedings of IEEE International Conference on Computer Vision, 2019. 3215–3223
- 44 Gou Y, Li B, Liu Z, et al. CLEARER: multi-scale neural architecture search for image restoration. In: Proceedings of International Conference on Neural Information Processing Systems, 2020, 33: 17129–17140
- 45 Bau D, Zhu J Y, Strobelt H, et al. GAN dissection: visualizing and understanding generative adversarial networks. In: Proceedings of International Conference on Learning Representations, 2019
- 46 Bau D, Zhu J Y, Wulff J, et al. Seeing what a GAN cannot generate. In: Proceedings of IEEE International Conference on Computer Vision, 2019. 4502–4511
- 47 Goetschalckx L, Andonian A, Oliva A, et al. GANalyze: toward visual definitions of cognitive image properties. In: Proceedings of IEEE International Conference on Computer Vision, 2019. 5744–5753
- 48 Härkönen E, Hertzmann A, Lehtinen J, et al. GANSpace: discovering interpretable GAN controls. In: Proceedings of International Conference on Neural Information Processing Systems, 2020
- 49 Suzuki R, Koyama M, Miyato T, et al. Spatially controllable image synthesis with internal representation collaging. 2018. ArXiv:1811.10153
- 50 Bau D, Strobelt H, Peebles W, et al. Semantic photo manipulation with a generative image prior. ACM Trans Graph, 2019, 38: 1–11
- 51 Tewari A, Elgharib M, Bharaj G, et al. StyleRig: rigging StyleGAN for 3D control over portrait images. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 6142–6151
- 52 Abdal R, Zhu P, Mitra N J, et al. StyleFlow: attribute-conditioned exploration of StyleGAN-generated images using conditional continuous normalizing flows. ACM Trans Graph, 2021, 40: 1–21
- 53 Menon S, Damian A, Hu S, et al. Pulse: self-supervised photo upsampling via latent space exploration of generative models. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 2437–2445
- 54 Richardson E, Alaluf Y, Patashnik O, et al. Encoding in style: a StyleGAN encoder for image-to-image translation. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 2287–2296
- 55 Chan K C, Wang X, Xu X, et al. GLEAN: generative latent bank for large-factor image super-resolution. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 14245–14254
- 56 Wang X, Li Y, Zhang H, et al. Towards real-world blind face restoration with generative facial prior. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 9168–9178
- 57 Yang T, Ren P, Xie X, et al. GAN prior embedded network for blind face restoration in the wild. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 672–681
- 58 LeCun Y, Bengio Y, Hinton G. Deep learning. Nature, 2015, 521: 436–444
- 59 Deng J, Dong W, Socher R, et al. ImageNet: a large-scale hierarchical image database. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2009. 248–255
- 60 Lee C H, Liu Z, Wu L, et al. MaskGAN: towards diverse and interactive facial image manipulation. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 5549–5558
- 61 Lecun Y, Bottou L, Bengio Y, et al. Gradient-based learning applied to document recognition. Proc IEEE, 1998, 86: 2278–2324
- 62 Netzer Y, Wang T, Coates A, et al. Reading digits in natural images with unsupervised feature learning. In: Proceedings

of NIPS Workshop on Deep Learning and Unsupervised Feature Learning, 2011

- 63 Krizhevsky A. Learning multiple layers of features from tiny images. 2009. https://www.cs.toronto.edu/ kriz/learning-features-2009-TR.pdf
- 64 Liu Z, Yan S, Luo P, et al. Fashion landmark detection in the wild. In: Proceedings of European Conference on Computer Vision, 2016. 229–245
- 65 Cordts M, Omran M, Ramos S, et al. The cityscapes dataset for semantic urban scene understanding. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2016. 3213–3223
- 66 Shao S, Li Z, Zhang T, et al. Objects365: a large-scale, high-quality dataset for object detection. In: Proceedings of IEEE International Conference on Computer Vision, 2019. 8430–8439
- 67 Zhou B, Lapedriza A, Khosla A, et al. Places: a 10 million image database for scene recognition. IEEE Trans Pattern Anal Mach Intell, 2017, 40: 1452–1464
- 68 Krasin I, Duerig T, Alldrin N, et al. OpenImages: a public dataset for large-scale multi-label and multi-class image classification. 2017. https://storage.googleapis.com/openimages/web/index.html
- 69 Salimans T, Goodfellow I, Zaremba W, et al. Improved techniques for training GANs. In: Proceedings of International Conference on Neural Information Processing Systems, 2016
- 70 Szegedy C, Vanhoucke V, Ioffe S, et al. Rethinking the inception architecture for computer vision. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2016. 2818–2826
- 71 Gurumurthy S, Sarvadevabhatla S R K, Babu R V. DeliGAN: generative adversarial networks for diverse and limited data. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2017. 166–174
- 72 Che T, Li Y, Jacob A P, et al. Mode regularized generative adversarial networks. In: Proceedings of International Conference on Learning Representations, 2017
- 73 Zhou Z, Zhang W, Wang J. Inception score, label smoothing, gradient vanishing and $-\log(D(x))$ alternative. 2017. ArXiv:1708.01729
- 74 Zhou Z, Cai H, Rong S, et al. Activation maximization generative adversarial nets. In: Proceedings of International Conference on Learning Representations, 2018
- 75 Heusel M, Ramsauer H, Unterthiner T, et al. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: Proceedings of International Conference on Neural Information Processing Systems, 2017
- 76 Bonneel N, Rabin J, Peyré G, et al. Sliced and radon Wasserstein barycenters of measures. J Math Imaging Vision, 2015, 51: 22–45
- 77 Kolouri S, Nadjahi K, Simsekli U, et al. Generalized sliced Wasserstein distances. 2019. ArXiv:1902.00434
- 78 Shmelkov K, Schmid C, Alahari K. How good is my GAN? In: Proceedings of European Conference on Computer Vision, 2018. 213–229
- 79 Kynkäänniemi T, Karras T, Laine S, et al. Improved precision and recall metric for assessing generative models. 2019. ArXiv:1904.06991
- 80 Khrulkov V, Oseledets I. Geometry score: a method for comparing generative adversarial networks. In: Proceedings of International Conference on Learning Representations, 2018. 2621–2629
- 81 Wang Z, Bovik A C, Sheikh H R, et al. Image quality assessment: from error visibility to structural similarity. IEEE Trans Image Process, 2004, 13: 600–612
- 82 Zhang R, Isola P, Efros A A, et al. The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2018. 586–595
- 83 Borji A. Pros and cons of GAN evaluation measures. Comput Vision Image Understanding, 2019, 179: 41–65
- 84 Wang Z, She Q, Ward T E. Generative adversarial networks in computer vision. ACM Comput Surv, 2022, 54: 1–38
- 85 Kang M, Shin J, Park J. StudioGAN: a taxonomy and benchmark of gans for image synthesis. 2022. ArXiv:2206.09479
- 86 Mescheder L, Geiger A, Nowozin S. Which training methods for GANs do actually converge? In: Proceedings of International Conference on Learning Representations, 2018. 3481–3490
- 87 Huang X, Belongie S. Arbitrary style transfer in real-time with adaptive instance normalization. In: Proceedings of IEEE International Conference on Computer Vision, 2017. 1501–1510
- 88 Tancik M, Srinivasan P, Mildenhall B, et al. Fourier features let networks learn high frequency functions in low dimensional domains. In: Proceedings of International Conference on Neural Information Processing Systems, 2020. 7537–7547
- 89 Mirza M, Osindero S. Conditional generative adversarial nets. 2014. ArXiv:1411.1784
- 90 Perarnau G, van de Weijer J, Raducanu B, et al. Invertible conditional GANs for image editing. In: Proceedings of NeurIPSW, 2016
- 91 Abdal R, Qin Y, Wonka P. Image2StyleGAN: how to embed images into the stylegan latent space? In: Proceedings of IEEE International Conference on Computer Vision, 2019. 4432–4441
- 92 Liu Y, Li Q, Sun Z, et al. Style intervention: how to achieve spatial disentanglement with style-based generators? 2020. ArXiv:2011.09699
- 93 Wu Z, Lischinski D, Shechtman E. Stylespace analysis: disentangled controls for StyleGAN image generation. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 12863–12872
- 94 Xu J, Xu H, Ni B, et al. Hierarchical style-based networks for motion synthesis. In: Proceedings of European Conference on Computer Vision, 2020. 178–194
- 95 Zhang L, Bai X, Gao Y. SalS-GAN: spatially-adaptive latent space in StyleGAN for real image embedding. In: Proceedings of ACM International Conference on Multimedia, 2021. 5176–5184
- 96 Zhu P, Abdal R, Qin Y, et al. Improved StyleGAN embedding: where are the good latents? 2020. ArXiv:2012.09036
- 97 Abdal R, Qin Y, Wonka P. Image2StyleGAN++: how to edit the embedded images? In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 8296–8305
- 98 Kang K, Kim S, Cho S. GAN inversion for out-of-range images with geometric transformations. In: Proceedings of IEEE International Conference on Computer Vision, 2021. 13941–13949
- 99 Cherepkov A, Voynov A, Babenko A. Navigating the GAN parameter space for semantic image editing. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 3671–3680
- 100 Feng Q, Shah V, Gadde R, et al. Near perfect GAN inversion. 2022. ArXiv:2202.11833
- 101 Donahue J, Krähenbühl P, Darrell T. Adversarial feature learning. In: Proceedings of International Conference on Learning Representations, 2017
- 102 Dumoulin V, Belghazi I, Poole B, et al. Adversarially learned inference. 2016. ArXiv:1606.00704
- 103 Zhu J Y, Krähenbühl P, Shechtman E, et al. Generative visual manipulation on the natural image manifold. In: Proceedings

of European Conference on Computer Vision, 2016. 597-613

- 104 Creswell A, Bharath A A. Inverting the generator of a generative adversarial network. IEEE Trans Neural Netw Learn Syst, 2019, 30: 1967–1974
- 105 Lipton Z C, Tripathi S. Precise recovery of latent vectors from generative adversarial networks. In: Proceedings of International Conference on Learning Representations Workshops, 2017
- 106 Shah V, Hegde C. Solving linear inverse problems using GAN priors: an algorithm with provable guarantees. In: Proceedings of International Conference on Acoustics, Speech and Signal Processing, 2018. 4609–4613
- 107 Ma F, Ayaz U, Karaman S. Invertibility of convolutional generative networks from partial measurements. In: Proceedings of International Conference on Neural Information Processing Systems, 2018. 9651–9660
- 108 Raj A, Li Y, Bresler Y. GAN-based projector for faster recovery with convergence guarantees in linear inverse problems. In: Proceedings of IEEE International Conference on Computer Vision, 2019. 5602–5611
- 109 Bau D, Zhu J Y, Wulff J, et al. Inverting layers of a large generator. In: Proceedings of International Conference on Learning Representations Workshops, 2019. 4
- 110 Shen Y, Gu J, Tang X, et al. Interpreting the latent space of GANs for semantic face editing. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 9243–9252
- 111 Daras G, Odena A, Zhang H, et al. Your local GAN: designing two dimensional local attention mechanisms for generative models. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 14531–14539
- 112 Gu J, Shen Y, Zhou B. Image processing using multi-code GAN prior. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 3012–3021
- 113 Anirudh R, Thiagarajan J J, Kailkhura B, et al. MimicGAN: robust projection onto image manifolds with corruption mimicking. Int J Comput Vis, 2020, 128: 2459–2477
- 114 Pan X, Zhan X, Dai B, et al. Exploiting deep generative prior for versatile image restoration and manipulation. IEEE Trans Pattern Anal Mach Intell, 2022, 44: 7474–7489
- 115 Viazovetskyi Y, Ivashkin V, Kashin E. StyleGAN2 distillation for feed-forward image manipulation. In: Proceedings of European Conference on Computer Vision, 2020. 170–186
- 116 Collins E, Bala R, Price B, et al. Editing in style: uncovering the local semantics of GANs. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 5771–5780
- 117 Pidhorskyi S, Adjeroh D A, Doretto G. Adversarial latent autoencoders. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 14104–14113
- 118 Huh M, Zhang R, Zhu J Y, et al. Transforming and projecting images into class-conditional generative networks. In: Proceedings of European Conference on Computer Vision, 2020. 17–34
- 119 Nitzan Y, Bermano A, Li Y, et al. Face identity disentanglement via latent space mapping. ACM Trans Graph, 2020, 39: 1-14
- 120 Aberdam A, Simon D, Elad M. When and how can deep generative models be inverted? 2020. ArXiv:2006.15555
- 121 Guan S, Tai Y, Ni B, et al. Collaborative learning for faster StyleGAN embedding. 2020. ArXiv:2007.01758
- 122 Shen Y, Zhou B. Closed-form factorization of latent semantics in GANs. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 1532–1540
- 123 Xu Y, Shen Y, Zhu J, et al. Generative hierarchical features from synthesizing images. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 4432–4442
- 124 Tewari A, Elgharib M, R M B, et al. PIE: portrait image embedding for semantic control. ACM Trans Graph, 2020, 39: 1-14
- 125 Bartz C, Bethge J, Yang H, et al. One model to reconstruct them all: a novel way to use the stochastic noise in StyleGAN. In: Proceedings of British Machine Vision Association, 2020
- 126 Wang H P, Yu N, Fritz M. Hijack-GAN: unintended-use of pretrained, black-box GANs. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 7872–7881
- 127 Zhuang P, Koyejo O O, Schwing A. Enjoy your editing: controllable GANs for image editing via latent space navigation.In: Proceedings of International Conference on Learning Representations, 2021
- 128 Alaluf Y, Patashnik O, Cohen-Or D. Only a matter of style: age transformation using a style-based regression model. ACM Trans Graph, 2021, 40: 1–12
- 129 Tov O, Alaluf Y, Nitzan Y, et al. Designing an encoder for StyleGAN image manipulation. ACM Trans Graph, 2021, 40: 1–14
- 130 Patashnik O, Wu Z, Shechtman E, et al. StyleCLIP: text-driven manipulation of StyleGAN imagery. In: Proceedings of IEEE International Conference on Computer Vision, 2021. 2085–2094
- 131 Chai L, Wulff J, Isola P. Using latent space regression to analyze and leverage compositionality in GANs. In: Proceedings of International Conference on Learning Representations, 2021
- 132 Chai L, Zhu J Y, Shechtman E, et al. Ensembling with deep generative views. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 14997–15007
- 133 Alaluf Y, Patashnik O, Cohen-Or D. ReStyle: a residual-based StyleGAN encoder via iterative refinement. In: Proceedings of IEEE International Conference on Computer Vision, 2021. 6711–6720
- 134 Wei T, Chen D, Zhou W, et al. E2Style: improve the efficiency and effectiveness of StyleGAN inversion. IEEE Trans Image Process, 2022, 31: 3267–3280
- 135 Xu Y, Du Y, Xiao W, et al. From continuity to editability: inverting GANs with consecutive images. In: Proceedings of IEEE International Conference on Computer Vision, 2021. 13910–13918
- 136 Wang T, Zhang Y, Fan Y, et al. High-fidelity GAN inversion for image attribute editing. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2022. 11379–11388
- 137 Schwettmann S, Hernandez E, Bau D, et al. Toward a visual concept vocabulary for GAN latent space. In: Proceedings of IEEE International Conference on Computer Vision, 2021. 6804–6812
- 138 Alaluf Y, Tov O, Mokady R, et al. HyperStyle: StyleGAN inversion with hypernetworks for real image editing. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2022. 18511–18521
- 139 Peebles W, Zhu J Y, Zhang R, et al. GAN-supervised dense visual alignment. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2022. 13470–13481
- 140 Dinh T M, Tran A T, Nguyen R, et al. HyperInverter: improving StyleGAN inversion via hypernetwork. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2022. 11389–11398
- 141 Alaluf Y, Patashnik O, Wu Z, et al. Third time's the charm? Image and video editing with StyleGAN3. 2022.

ArXiv:2201.13433

- 142 Frühstück A, Singh K K, Shechtman E, et al. InsetGAN for full-body image generation. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2022. 7723–7732
- 143 Wu Y, Yang Y L, Jin X. HairMapper: removing hair from portraits using gans. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2022. 4227–4236
- 144 Parmar G, Li Y, Lu J, et al. Spatially-adaptive multilayer selection for GAN inversion and editing. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2022. 11399–11409
- 145 Zhou B, Zhao H, Puig X, et al. Scene parsing through ADE20K dataset. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2017. 633–641
- 146 Chen B C, Chen C S, Hsu W H. Cross-age reference coding for age-invariant face recognition and retrieval. In: Proceedings of European Conference on Computer Vision, 2014. 768–783
- 147 Lin T Y, Maire M, Belongie S, et al. Microsoft COCO: common objects in context. In: Proceedings of European Conference on Computer Vision, 2014. 740–755
- 148 Wah C, Branson S, Welinder P, et al. The Caltech-UCSD birds-200-2011 dataset. 2011. http://www.vision.caltech.edu/visipedia/CUB-200.html
- 149 Anonymous, The Danbooru Community, Branwen G. Danbooru2021: a large-scale crowdsourced and tagged anime illustration dataset. 2021. https://www.gwern.net/Danbooru
- 150 Nilsback M E, Zisserman A. Automated flower classification over a large number of classes. In: Proceedings of the 6th Indian Conference on Computer Vision, Graphics & Image Processing, 2008. 722–729
- 151 Huang G B, Mattar M, Berg T, et al. Labeled faces in the wild: a database forstudying face recognition in unconstrained environments. In: Proceedings of Workshop on Faces in 'Real-Life' Images: Detection, Alignment, and Recognition, 2008
- 152 Skorokhodov I, Sotnikov G, Elhoseiny M. Aligning latent and image spaces to connect the unconnectable. In: Proceedings of IEEE International Conference on Computer Vision, 2021. 14144–14153
- 153 Lake B M, Salakhutdinov R, Tenenbaum J B. Human-level concept learning through probabilistic program induction. Science, 2015, 350: 1332–1338
- 154 Zhou B, Lapedriza A, Xiao J, et al. Learning deep features for scene recognition using places database. In: Proceedings of International Conference on Neural Information Processing Systems, 2014
- 155 Parkhi O M, Vedaldi A, Zisserman A, et al. Cats and dogs. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2012. 3498–3505
- 156 Livingstone S R, Russo F A. The Ryerson Audio-Visual Database of Emotional Speech and Song (RAVDESS): a dynamic, multimodal set of facial and vocal expressions in North American English. Plos One, 2018, 13: e0196391
- 157 Krause J, Stark M, Deng J, et al. 3D object representations for fine-grained categorization. In: Proceedings of IEEE International Conference on Computer Vision Workshops, 2013. 554–561
- 158 Naik N, Philipoom J, Raskar R, et al. Streetscore-predicting the perceived safety of one million streetscapes. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition Workshops, 2014. 779–785
- 159 Laffont P Y, Ren Z, Tao X, et al. Transient attributes for high-level understanding and editing of outdoor scenes. ACM Trans Graph, 2014, 33: 1–11
- 160 Yu A, Grauman K. Fine-grained visual comparisons with local learning. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2014. 192–199
- 161 Liu D C, Nocedal J. On the limited memory BFGS method for large scale optimization. Math Programming, 1989, 45: 503–528
- 162 Kingma D P, Ba J. Adam: a method for stochastic optimization. 2014. ArXiv:1412.6980
- 163 Deng J, Guo J, Xue N, et al. ArcFace: additive angular margin loss for deep face recognition. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2019. 4690–4699
- 164 Huang Y, Wang Y, Tai Y, et al. CurricularFace: adaptive curriculum learning loss for deep face recognition. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 5901–5910
- 165 He K, Fan H, Wu Y, et al. Momentum contrast for unsupervised visual representation learning. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2020. 9729–9738
- 166 Donahue J, Simonyan K. Large scale adversarial representation learning. In: Proceedings of International Conference on Neural Information Processing Systems, 2019. 32
- 167 Kingma D P, Dhariwal P. Glow: generative flow with invertible 1× 1 convolutions. In: Proceedings of International Conference on Neural Information Processing Systems, 2018. 10236–10245
- 168 Ho J, Jain A, Abbeel P. Denoising diffusion probabilistic models. In: Proceedings of International Conference on Neural Information Processing Systems, 2020. 6840–6851
- 169 Tousi A, Jeong H, Han J, et al. Automatic correction of internal units in generative neural networks. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 7932–7940
- 170 Bau D, Zhou B, Khosla A, et al. Network dissection: quantifying interpretability of deep visual representations. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2017. 6541–6549
- 171 Carter S, Armstrong Z, Schubert L, et al. Activation atlas. Distill, 2019. Https://distill.pub/2019/activation-atlas
- 172 Bau D, Liu S, Wang T, et al. Rewriting a deep generative model. In: Proceedings of European Conference on Computer Vision, 2020. 351–369
- 173 Langner O, Dotsch R, Bijlstra G, et al. Presentation and validation of the Radboud Faces Database. Cognition Emotion, 2010, 24: 1377–1388
- 174 Ramesh A, Choi Y, LeCun Y. A spectral regularizer for unsupervised disentanglement. 2018. ArXiv:1812.01161
- 175 Chen X, Duan Y, Houthooft R, et al. InfoGAN: interpretable representation learning by information maximizing generative adversarial nets. In: Proceedings of International Conference on Neural Information Processing Systems, 2016. 2172–2180
- 176 Peebles W, Peebles J, Zhu J Y, et al. The hessian penalty: a weak prior for unsupervised disentanglement. In: Proceedings of European Conference on Computer Vision, 2020. 581–597
- 177 Zhu X, Xu C, Tao D. Learning disentangled representations with latent variation predictability. In: Proceedings of European Conference on Computer Vision, 2020. 684–700
- 178 Zhu X, Xu C, Tao D. Where and what? Examining interpretable disentangled representations. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 5861–5870
- 179 Wei Y, Shi Y, Liu X, et al. Orthogonal jacobian regularization for unsupervised disentanglement in image generation. In: Proceedings of IEEE International Conference on Computer Vision, 2021. 6721–6730

- 180 He Z, Kan M, Shan S. EigenGAN: layer-wise eigen-learning for GANs. In: Proceedings of IEEE International Conference on Computer Vision, 2021. 14408–14417
- 181 Jahanian A, Chai L, Isola P. On the "steerability" of generative adversarial networks. In: Proceedings of International Conference on Learning Representations, 2020
- 182 Zhu J, Shen Y, Xu Y, et al. Region-based semantic factorization in GANs. 2022. ArXiv:2202.09649
- 183 Wang B, Ponce C R. A geometric analysis of deep generative image models and its applications. In: Proceedings of International Conference on Learning Representations, 2021
- 184 Tzelepis C, Tzimiropoulos G, Patras I. WarpedGANSpace: finding non-linear RBF paths in GAN latent space. In: Proceedings of IEEE International Conference on Computer Vision, 2021. 6393–6402
- 185 Wang X, Yu K, Dong C, et al. Deep network interpolation for continuous imagery effect transition. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2019. 1692–1701
- 186 Selvaraju R R, Cogswell M, Das A, et al. Grad-CAM: visual explanations from deep networks via gradient-based localization. In: Proceedings of IEEE International Conference on Computer Vision, 2017. 618–626
- 187 Pan X, Dai B, Liu Z, et al. Do 2D GANs know 3D shape? Unsupervised 3D shape reconstruction from 2D image GANs. 2020. ArXiv:2011.00844
- 188 Zhang J, Chen X, Cai Z, et al. Unsupervised 3D shape completion through GAN inversion. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 1768–1777
- 189 Kingma D P, Welling M. Auto-encoding variational Bayes. 2013. ArXiv:1312.6114
- 190 van den Oord A, Kalchbrenner N, Kavukcuoglu K. Pixel recurrent neural networks. In: Proceedings of International Conference on Machine Learning, 2016. 1747–1756
- 191 Ramesh A, Dhariwal P, Nichol A, et al. Hierarchical text-conditional image generation with clip latents. 2022. ArXiv:2204.06125
- 192 Saharia C, Chan W, Saxena S, et al. Photorealistic text-to-image diffusion models with deep language understanding. 2022. ArXiv:2205.11487
- 193 Zhang D, Han J, Cheng G, et al. Weakly supervised object localization and detection: a survey. IEEE Trans Pattern Anal Mach Intell, 2021, 44: 5866–5885
- 194 Han J, Zhang D, Cheng G, et al. Advanced deep-learning techniques for salient and category-specific object detection: a survey. IEEE Signal Process Mag, 2018, 35: 84–100
- 195 Zhang D, Tian H, Han J. Few-cost salient object detection with adversarial-paced learning. In: Proceedings of International Conference on Neural Information Processing Systems, 2020. 33: 12236–12247
- 196 Frid-Adar M, Klang E, Amitai M, et al. Synthetic data augmentation using GAN for improved liver lesion classification. In: Proceedings of International Symposium on Biomedical Imaging, 2018. 289–293
- 197 Huang S W, Lin C T, Chen S P, et al. AugGAN: cross domain adaptation with gan-based data augmentation. In: Proceedings of European Conference on Computer Vision, 2018. 718–731
- 198 Zhang Y, Ling H, Gao J, et al. DatasetGAN: efficient labeled data factory with minimal human effort. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 10145–10155
- 199 Han M, Zheng H, Wang C, et al. Leveraging GAN priors for few-shot part segmentation. In: Proceedings of ACM International Conference on Multimedia, 2022. 1339–1347
- 200 Schlegl T, Seeböck P, Waldstein S M, et al. f-AnoGAN: fast unsupervised anomaly detection with generative adversarial networks. Med Image Anal, 2019, 54: 30–44
- 201 Dunn I, Pouget H, Melham T, et al. Adaptive generation of unrestricted adversarial inputs. 2019. ArXiv:1905.02463
- 202 Wang X, He K, Hopcroft J E. At-GAN: a generative attack model for adversarial transferring on generative adversarial nets. 2019. ArXiv:1904.07793
- 203 Ojha U, Li Y, Lu J, et al. Few-shot image generation via cross-domain correspondence. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2021. 10743–10752
- 204 Gu J, Liu L, Wang P, et al. StyleNeRF: a style-based 3D aware generator for high-resolution image synthesis. In: Proceedings of International Conference on Learning Representations, 2022
- 205 He J, Shi W, Chen K, et al. GCFSR: a generative and controllable face super resolution method without facial and GAN priors. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2022. 1889–1898
- 206 Li X, Chen C, Lin X, et al. From face to natural image: learning real degradation for blind image super-resolution. In: Proceedings of European Conference on Computer Vision, 2022
- 207 Li B, Liu X, Hu P, et al. All-in-one image restoration for unknown corruption. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 2022. 17452–17462
- 208 Lyu Z, Xu X, Yang C, et al. Accelerating diffusion models via early stop of the diffusion process. 2022. ArXiv:2205.12524
- 209 Grover A, Dhar M, Ermon S. Flow-GAN: combining maximum likelihood and adversarial learning in generative models. In: Proceedings of AAAI Conference on Artificial Intelligence, 2018