



Optimizing and interpreting the latent space of the conditional text-to-image GANs

Zhenxing Zhang¹ · Lambert Schomaker¹

Received: 27 March 2023 / Accepted: 20 October 2023 / Published online: 21 November 2023
© The Author(s) 2023

Abstract

Text-to-image generation intends to automatically produce a photo-realistic image, conditioned on a textual description. To facilitate the real-world applications of text-to-image synthesis, we focus on studying the following three issues: (1) How to ensure that generated samples are believable, realistic or natural? (2) How to exploit the latent space of the generator to edit a synthesized image? (3) How to improve the explainability of a text-to-image generation framework? We introduce two new data sets for benchmarking, i.e., the *Good & Bad*, bird and face, data sets consisting of successful as well as unsuccessful generated samples. This data set can be used to effectively and efficiently acquire high-quality images by increasing the probability of generating *Good* latent codes with a separate, new classifier. Additionally, we present a novel algorithm which identifies semantically understandable directions in the latent space of a conditional text-to-image GAN architecture by performing independent component analysis on the pre-trained weight values of the generator. Furthermore, we develop a background-flattening loss (BFL), to improve the background appearance in the generated images. Subsequently, we introduce linear-interpolation analysis between pairs of text keywords. This is extended into a similar triangular ‘linguistic’ interpolation. The visual array of interpolation results gives users a deep look into what the text-to-image synthesis model has learned within the linguistic embeddings. Experimental results on the recent DiverGAN generator, pre-trained on three common benchmark data sets demonstrate that our classifier achieves a better than 98% accuracy in predicting Good/Bad classes for synthetic samples and our proposed approach is able to derive various interpretable semantic properties for the text-to-image GAN model.

Keywords Generative adversarial network (GAN) · Text-to-image generation · A *Good & Bad* data set · Latent-space manipulation

1 Introduction

The task of text-to-image synthesis aims at automatically generating high-quality and semantically consistent images, given natural-language descriptions. It has recently gathered increasing interest from researchers due to its numerous potential applications, e.g., data augmentation for training image classifiers, photoediting according to textual descriptions, the education of young children, etc.

With the advances in the generative adversarial network (GAN) and the conditional generative adversarial network (cGAN) [1], text-to-image generation has achieved promising progress in both image quality and semantic consistency. Nevertheless, it remains extremely challenging to coerce a conditional text-to-image GAN model to generate, with high probability, believable and natural images.¹ Here, we want to address three research questions in order to improve easier methods to generate high-quality images:

1. How to ensure that generated samples are believable, realistic or natural?
2. How to exploit the latent space of the generator to edit a synthesized image?

✉ Zhenxing Zhang
z.zhang@rug.nl

Lambert Schomaker
l.r.b.schomaker@rug.nl

¹ Bernoulli Institute, University of Groningen, Groningen, The Netherlands

¹ A preliminary version of this paper was submitted to ArXiv on [2].

3. How to improve the explainability of a text-to-image generation framework?

One particular disadvantage of synthetic image-generation algorithms is that the performance evaluation is more difficult than is the case in classification problems where a ‘hard’ accuracy can be computed. In case of the cGAN this issue is most clearly present for end users: *How to ensure that generated images are believable, realistic or natural?* In current literature, the good examples are often cherry picked while occasionally also the less successful samples are shown. However, for actual use in data augmentation or in artistic applications, one would like to guarantee that generated images are good, i.e., of a sufficiently believable natural quality. Given the high dimensionality of latent codes, there is a very high prior probability of non-successful patterns to be generated for a given input noise probe. How to construct a random latent-code generator with an increased probability of drawing successful samples? After the generator/discriminator pair has done its best effort, apparently additional constraints are necessary.

Here, we intend to train a classifier to accurately distinguish successful synthesized samples from unsuccessful generated pictures after training a text-to-image generation framework. This is based on the assumption that there is a nonlinear boundary separating high-resolution images from inadequate samples in the fake image space. To this end, we created a *Good & Bad* data set, both for a bird and a face-image collection, which consists of a large number of realistic as well as implausible samples synthesized by the recent DiverGAN [3] that was pretrained on the CUB bird [4] data set and the Multi-Modal CelebA-HQ data set [5], respectively. We choose these samples by following strict principles in order to ensure the quality of the selected images. To acquire a superior classifier, we train the CNN model (e.g., ResNet [6]) from the pre-trained weights on our *Good & Bad* data set. We expect that the well-trained network can correctly predict the quality class of synthesized images. Therefore, we are able to effectively and efficiently derive photo-realistic images from the synthesized samples while obtaining corresponding *Good* latent vectors. More importantly, the discovery of *Good* latent codes provides a strong basis for further research, such as data augmentation and latent-space manipulation.

Latent vectors contributing to diversity play a significant role in the image-generation process. Recent works [7–9] reveal that there exists a wide range of meaningful semantic factors in the latent space of a GAN, such as facial attributes and head poses for face synthesis and layout for scene generation [9]. These semantically understandable control directions can be utilized for disentangled image editing, like semantic face editing and scene manipulation. By moving the latent code of a

synthetic sample toward and backward the direction, we are able to vary the desired attribute while keeping other image contents unchanged. That is to say, given a successful latent code, we can derive a wealth of similar but semantically diverse pleasing images via latent-space navigation. To better facilitate the application of text-to-image synthesis, we need to address the question: *How to identify useful control directions in the latent space of a conditional text-to-image GAN model?* While current approaches mainly focus on studying the latent space of a GAN, there still is a lack of understanding of the relationship between the latent space of a cGAN and the explainable semantic space in which a synthetic sample is embedded.

In this paper, we present a novel algorithm to capture the interpretable latent-space semantic properties for a text-to-image synthesis model. Considering the fact that identified directions denote different semantic factors of the edited object (e.g., pose and smile for the face model), we argue that these vectors should be fully independent rather than just uncorrelated. Based on recent studies [7, 8], we assume that the pre-trained weights of a conditional text-to-image GAN architecture contain a set of useful directions. In fact, the initial linear layer projects the latent vector to the visual feature map, where a latent space is transformed into another space and ultimately into an output image. To acquire both independent and orthogonal components, we introduce the independent component analysis (ICA) algorithm under an additional orthogonality constraint [10] to investigate the pre-trained weight matrix of the first dense layer. In addition, we mathematically show that Semantic Factorization (SeFa) [7], GANSpace [8] and regular PCA typically achieve almost the identical results when sampling enough data for GANSpace. Furthermore, we develop a *Background-Flattening Loss* (BFL), to improve the background appearance in the edited sample. Multiple interesting latent-space directions found by our presented algorithm are visualized in Fig. 1.

We expect that our proposed semantic-discovery method can provide valuable insight into the correlation between latent vectors and image variations. However, it remains particularly difficult to explain what a conditional text-to-image GAN model has learned within the text space. *How to understand the relation between the textual (linguistic) probes and the generated image factors?* This constitutes the last research topic of the current contribution. To alleviate the problem, we qualitatively analyze the roles played by the linguistic embeddings in the generated-image semantic space through linear interpolation analysis between pairs of keywords. We show that although semantic properties contained in the picture change continuously in the latent space, the appearance of the image does not always vary smoothly along with the contrasting

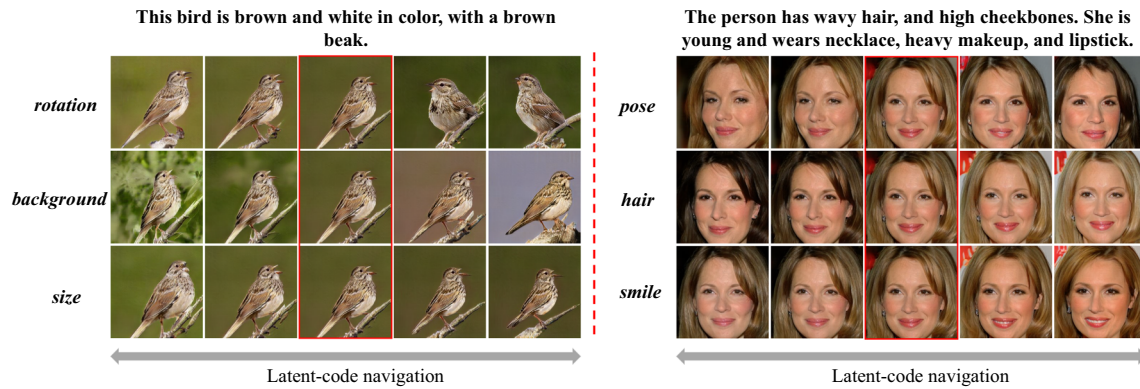


Fig. 1 Interpretable latent-space directions identified in DiverGAN [3] that was pre-trained on the CUB bird [4] (left side) and Multi-Modal CelebA-HQ [5] (right side) data sets. For each set of pictures, the middle column is the original image based on a *Good* latent code,

while the samples on the left and right of it are the output by freezing the textual description and moving the latent vector backward and forward from the center, over the axis discovered by our proposed algorithm

word embeddings. In addition, we extend a pairwise linear interpolation to a triangular interpolation for simultaneously investigating three keywords in the given textual description.

The recent DiverGAN [3] has the ability to adopt a generator/discriminator pair to synthesize diverse and high-quality samples, given a textual description and different injected noise on the latent vector. We therefore carry out a series of experiments on the DiverGAN generator that was trained on three popular text-to-image data sets (i.e., the CUB bird [4], MS COCO [11] and Multi-Modal CelebA-HQ [5] data sets). The experimental results in the current study represent an improvement in performance and explainability in the analyzed algorithm [3]. Meanwhile, our well-trained classifier achieves impressive classification accuracy (bird: 98.09% and face: 99.16%) on the *Good* & *Bad* data set and our proposed semantic-discovery algorithm can lead to a more precise control over the latent space of the DiverGAN model, which validate the effectiveness of our presented methods. The contributions of this work can be summarized as follows:

- We construct two new *Good* & *Bad* data sets to study how to ensure that generated images are believable while training two corresponding classifiers to separate successful generated images from unsuccessful synthetic samples.
- We introduce the ICA algorithm to identify meaningful attributes in the latent space of a conditional text-to-image GAN model. Simultaneously, we analyze the correspondences between SeFa, GANSpace and regular PCA.
- We introduce linear interpolation analysis between pairs of contrastive keywords and a similar triangular ‘linguistic’ interpolation for an improved explainability of a text-to-image generation architecture.

The remainder of the paper is organized as follows. We introduce the related works in Sect. 2. Section 3 briefly depicts the single-stage text-to-image framework and the corresponding latent space. In Sect. 4, we describe our proposed approach in detail. The experimental results are presented in Sect. 5, and Sect. 6 draws the conclusions.

2 Related works

In this section, we depict the research fields associated with our work, i.e., a GAN, cGAN-based text-to-image generation and latent-space manipulation.

2.1 Generative adversarial network (GAN)

Goodfellow et al. [12] presented the GAN paradigm that serves as a basic model for synthetic tasks via adversarial training and consists of a generator and a discriminator. A GAN has achieved state-of-the-art performance in a variety of applications, e.g., text-to-image synthesis [13], person image generation [14], face photo-sketch synthesis [15], image inpainting [16] and image de-raining [17], since it is capable of producing photo-realistic images.

The initial generator network of a GAN mainly comprises multi-layer perceptrons and rectifier linear activations, while the discriminator net utilizes the maxout network [18]. This type of architecture shows competitive samples with other generative models on simple image data sets, such as MNIST [19]. Moreover, researchers explore different structures of a GAN in order to further improve image quality. Denton et al. [20] designed a Laplacian pyramid framework of an adversarial network, namely LAPGAN that produces plausible pictures in a coarse-to-fine manner. Radford et al. [21] introduced a deep convolutional GAN (DCGAN) integrating convolutional layers

and Batch Normalization (BN) [22] into both a generator and a discriminator. Mirza et al. [1] proposed a cGAN by imposing conditional constraints (e.g., class labels, text descriptions and low-resolution images) on both the generator network and the discriminator network to obtain specific samples.

Recently, several models with a high-computational cost are introduced to yield visually plausible pictures. Zhang et al. [23] presented SAGAN which applies the self-attention mechanism [24] to effectively capture the semantic affinities between widely separated image regions. Brock et al. [25] developed a large-scale architecture based on SAGAN while deploying orthogonal regularization to the generator net, obtaining excellent performance on image diversity. Karras et al. [26] proposed a novel generator framework named StyleGAN where adaptive instance normalization is utilized to control the generator network. Kumar et al. [27] proposed to combine Phase-space reconstruction (PSR) with GANs to improve the prediction accuracy of stock price movement direction. The experimental results showed that the presented approach reduced the root-mean-square error value from 0.0585 to 0.0295 and the average processing time from 3 min 26 s to 2 min 8 s, validating its effectiveness. Jameel et al. [28] utilized a cGAN to produce a variety of new medical pictures for the purpose of data augmentation. They discovered that data augmentation using the cGAN performs better than traditional data-augmentation techniques. The current paper focuses on studying and improving the conditional text-to-image GAN approach.

2.2 cGAN in text-to-image generation

Owing to the success of a GAN on image quality, the task of text-to-image synthesis has achieved significant advances over the past few years. Existing approaches for text-to-image generation can be roughly cast into two categories: (1) multi-stage models and (2) single-stage methods.

Multi-stage models Zhang et al. [29] introduced a multi-stage architecture called StackGAN, in which each stage comprises a generator and a discriminator, and the generator of the next stage receives the result of the previous stage as the input. It is worth mentioning that StackGAN served as a solid basis for the future study of text-to-image synthesis. Xu et al. [30] proposed AttnGAN inserting a spatial attention module into the multi-stage framework to bridge the semantic gap between the words in a textual description and the related image subregions. The introduced spatial-attention mechanism is able to derive the relationship between the image subregions and the words in a sentence. The most relevant subregions to the words were particularly focused. Qiao et al. [31] presented MirrorGAN where an image-to-text model is leveraged to

guarantee the semantic consistency between natural-language descriptions and visual contents. Zhu et al. [32] designed DMGAN which introduce a dynamic-memory module to produce high-quality samples in the initial stage. Specifically, a memory writing gate and a response gate are designed to better fuse the textual and image information. OP-GAN presented by Hinz et al. [33] explicitly modeled the objects of an image while developing a new evaluation metric termed as semantic object accuracy. Cheng et al. [34] presented RiFeGAN, in which an attention-based caption-matching algorithm is utilized to enrich the textual description and a multi-captions attentional GAN model is exploited to yield photo-realistic pictures. It should be noted that the multi-stage models are unattractive due to complicated design and training.

Single-stage methods are an attempt to remedy this disadvantage. Reed et al. [35] were the first to attempt to employ the cGAN to synthesize specific images based on the given text descriptions. Tao et al. [36] proposed DFGAN where a matching-aware zero-centered gradient penalty loss is introduced to help stabilize the training of the conditional text-to-image GAN model. Z. Zhang et al. [37] designed DTGAN by utilizing spatial and channel attention modules and the conditional normalization to yield photo-realistic samples with a generator/discriminator pair. More importantly, they spread the sentence-related attention over many, even all layers of the generator networks. This allows for an influence of the text over features at various hierarchical levels in the pipeline architecture, from crude early features to abstract late features. H. Zhang et al. [38] developed XMC-GAN which studied contrastive learning in the context of text-to-image generation while producing visually plausible images via a simple single-stage framework. Zhang et al. [3] presented an efficient and effective single-stage framework called DiverGAN which is capable of generating diverse, plausible and semantically consistent images according to a natural-language description. To be specific, they proposed two novel word-level attention modules, i.e., a channel-attention module (CAM) and a pixel-attention module (PAM), which allow DiverGAN to effectively disentangle the attributes of the text description while accurately controlling the regions of synthetic images. Xia et al. [5] developed TediGAN for image synthesis and control with natural-language descriptions. TediGAN comprises an inversion module, visual-language learning and instance-level optimization. Wang et al. [39] proposed a unified architecture for both text-to-image generation and text-driven image manipulation. In the current study we will adopt the DiverGAN generator due to its excellent performance on image quality and diversity in order to perform comprehensive experiments and facilitate the generation of good images.

2.3 Studies on the latent space of a GAN

Recent studies [7–9] on a GAN reveal that a latent space possesses a range of semantically understandable information (e.g., pose and smile for the face data set), which plays a vital role in detangled sample manipulation. We are able to realistically edit the generated image by moving its latent vector toward the direction corresponding to the desired attribute. Several methods have been proposed to capture interpretable semantic factors and mainly fall into two types: (1) unsupervised models and (2) supervised approaches.

Supervised latent-space manipulation Shen et al. [9] developed a framework termed as InterfaceGAN where labeled samples (e.g., gender and age) are utilized to train a linear Support Vector Machine (SVM) [40] and the acquired SVM boundaries lead to the meaningful manipulation of the facial attributes. Goetschalckx et al. [41] proposed GANalyze applying an accessor module to optimize the training process while learning the latent-space directions as the desired cognitive semantics.

Unsupervised latent-space manipulation Voynov et al. [42] introduced a matrix and a classifier to identify interpretable latent-space directions in an unsupervised fashion. Jahanian et al. [43] studied the attributes concerning color transformations and camera movements by operating source pictures. Härkönen et al. [8] designed a novel pipeline named GANSpace, which performed PCA on a series of collected latent vectors and employed obtained principal components as the meaningful directions in the latent space. Peebles et al. [44] presented the Hessian Penalty, a regularization term for the unsupervised discovery of useful semantic factors. Wang et al. [44] developed Hijack-GAN introducing an iterative scheme to control the image-generation process. Shen et al. [7] proposed Semantic Factorization (SeFa) which directly decomposed the weight matrix of a well-trained GAN model for semantic image editing. Patashnik et al. [45] introduced Contrastive Language-Image Pre-training (CLIP) models [46] to control the latent space using a text prompt. The current work aims at identifying controllable directions in the latent space of conditional text-to-image GAN models.

3 Preliminary

In this section, we briefly describe the single-stage text-to-image synthesis architecture and the corresponding latent space to help understand the issues we attempt to address.

Single-stage pipeline The single-stage text-to-image generation framework (illustrated in Fig. 2) is composed of

a generator network and a discriminator net, which are perceived as playing a minmax zero-sum game. Let $S = \{(C_i, I_i)\}_{i=1}^N$ denote a collection of N text-image pairs for training, where I_i is a picture and $C_i = (c_i^1, c_i^2, \dots, c_i^K)$ comprises K textual descriptions. Word-embedding vectors w and a sentence-embedding vector s are commonly acquired by applying a bidirectional Long Short-Term Memory (LSTM) network [47] on a natural-language description c_i randomly picked from C_i . After that, the generator $G(z, (w, s))$ is trained to produce a perceptually realistic and semantically related image \hat{I}_i according to a latent code z randomly sampled from a frozen distribution and word/sentence embedding vectors (w, s) . To be specific, $G(z, (w, s))$ consists of multiple layers where the first layer F_0 maps a latent code into a feature map and intermediate blocks typically leverage modulation modules (e.g., attention models [48, 49]) to reinforce the visual feature map to ensure image quality and semantic consistency. The last layer G_c transforms the feature map into the ultimate sample. Mathematically,

$$h_0 = F_0(z) \quad (1)$$

$$h_1 = B_1(h_0, (w, s)) \quad (2)$$

$$h_i = B_i(h_{i-1} \uparrow, (w, s)) \quad \text{for } i = 2, 3, \dots, 7 \quad (3)$$

$$\hat{I} = G_c(h_7) \quad (4)$$

where F_0 denotes a fully connected layer and B_i is a modulation block that facilitates the feature map with textual features.

Compared with $G(z, (w, s))$, the discriminator of the single-stage pipeline aims at distinguishing the real text-image pair (c_i, I_i) from the fake text-image pair (c_i, \hat{I}_i) .

Latent-space analysis For a pre-trained and fixed generator $G(z, (w, s))$, the quality of the generated sample depends on the random latent code z , word embeddings w and the corresponding sentence vector s . Consequently, the output of the network only relies on z when determining the input text description. It implicitly means that if we ignore the linguistic space of the conditional input-text probes, $G(z, (w, s))$ can be regarded as a deterministic function $G: \mathcal{Z} \rightarrow \mathcal{X}$. Here, \mathcal{Z} represents the latent space, in which the latent code $z \in R^l$ is commonly sampled from a 1-dimension Gaussian distribution. \mathcal{X} denotes the synthetic image space including visually realistic samples as well as implausible generated pictures. Moreover, the map from \mathcal{Z} to \mathcal{X} is not surjective [50]. Accordingly, even a superior text-to-image generation generator fails to ensure the quality of a synthesized sample, given random latent vectors. In order to promote the applicability of text-to-image generation in practice, this paper intends to optimize the latent space of a conditional text-to-image GAN model to

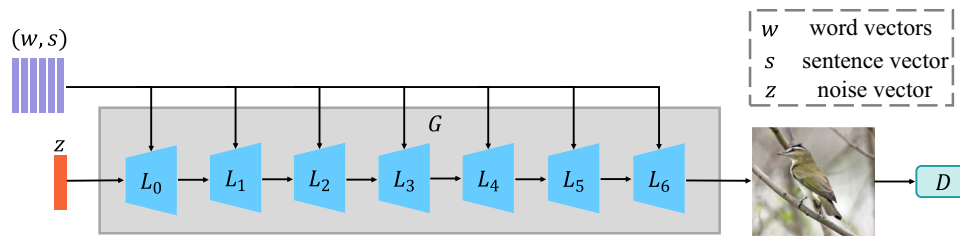


Fig. 2 A simplified single-stage text-to-image generation architecture consisting of a generator G and a discriminator D . The input of the generator is a random latent code z and the word/sentence embeddings (w, s) , and the output is a synthetic sample

effectively avoid unsuccessful synthetic samples while automatically obtaining high-quality images.

Latent-space manipulation It has been widely observed that the latent space of a GAN incorporates certain semantic information, like pose and size for the CUB bird data set. Suppose we have a *Good* latent code z_g that contributes to a successful generated sample, and a well-trained generator $G(z, (w, s))$ that can yield dissimilar and semantically consistent pictures according to different textual descriptions and injected noise, we target to manipulate the semantic factor of the successful synthesized sample via latent-space navigation. To this end, we need to first identify a series of semantically interpretable latent-space directions $N = (n_1, n_2, \dots, n_k)$, where $n_i \in R^l$ for all $i \in 1, 2, \dots, k$. Then, the attribute of the high-quality sample generated by z_g can be varied by editing z_g with $z_{ge} = z_g + \alpha n$, where α denotes the mix factor and $n \in R^l$ is the direction corresponding to the desired property.

4 Proposed methodology

This paper is part of a larger research, which focuses on proposing novel architectures for text-to-image generation while addressing significant issues present in a conditional text-to-image GAN model. An overall framework for the research is visualized in Fig. 3, consisting of DTGAN [37], DiverGAN [3] and three different research questions, i.e., RQ1: increasing the probability of generating natural images and RQ2: improving the (visual) explainability of a conditional text-to-image GAN model latent-space and RQ3: addressing textual explainability using the linguistic latent space point. Both DTGAN and DiverGAN are capable of adopting a single generator/discriminator pair to produce photo-realistic and semantically correlated image samples on the basis of given natural-language descriptions. In DTGAN, we presented dual-attention models, conditional adaptive instance-layer normalization and a new type of visual loss. In DiverGAN, we extended the sentence-level attention models introduced in DTGAN to word-level attention modules, in order to better control an

image-generation process using word features. By inserting a dense layer into the pipeline we were able to address the lack-of-diversity problem present in many current single-stage text-to-image GAN models. The explanation is that unlike feature maps, dense layers prevent the simple propagation of image fragments toward the output and force the network to fundamentally reorganize the visual features.

In this section, we elaborate on the proposed procedure automatically finding successful synthetic samples from generated images while acquiring corresponding *Good* latent codes. After that, based on a *Good* latent vector, we describe the independent component analysis method that identifies meaningful latent-space directions for a conditional text-to-image GAN model. Subsequently, we introduce linear interpolation analysis between contrastive keywords as well as a similar triangular ‘linguistic’ interpolation for an improved explainability of a text-to-image generation framework.

4.1 Discovering successful synthesized samples and *Good* latent codes

Given a fixed conditional text-to-image GAN model, the generator maps the latent space and the linguistic embeddings to the fake data distribution. It is well known that the synthetic sample space consists of high-resolution pictures from *Good* latent vectors as well as unreasonable images from *Bad* latent codes. However, we only need successful generated samples and corresponding *Good* latent codes for wide real-world applications. In this subsection, we concentrate on proposing a framework for recognizing plausible images from numerous synthesized samples while deriving corresponding *Good* latent vectors.

4.1.1 Pairwise linear interpolation of latent codes

It has been extensively observed [3, 9] that when performing the linear interpolation between a successful starting-point latent vector and a successful end-point latent code, the appearance and the semantics of generated samples change continuously. In addition, DiverGAN [3]

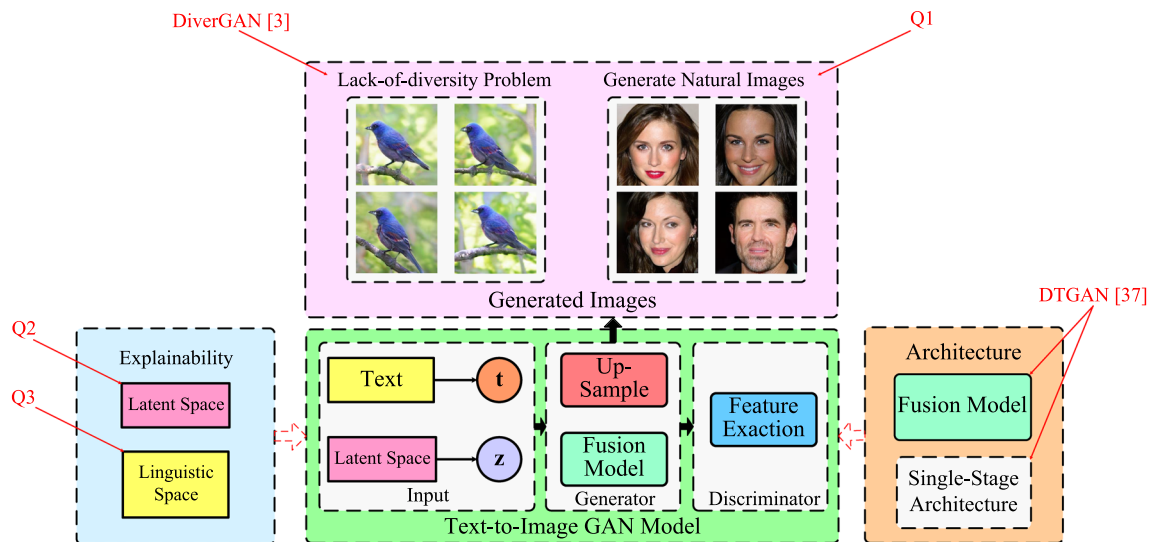


Fig. 3 The overall framework for the research questions of this paper. t and z denote the textual features of an input natural-language description and a latent code in the latent space, respectively

discovers that the generator is likely to synthesize a set of high-resolution pictures based on the pairwise linear interpolation between two *Good* latent codes. This would imply that there may be close relation between successful synthesized images in the fake data space. That is to say, we may acquire a range of visually realistic pictures by sampling the latent vectors around a *Good* latent code.

To further explore the semantic relationship between a plausible sample and an inadequate image in the synthetic image space, we visualize the samples generated by linearly interpolating a successful starting-point latent vector z_0 and an unsuccessful end-point latent code z_1 . To be specific, the pairwise linear interpolation of latent codes is defined as:

$$f(\gamma) = G((1 - \gamma)z_0 + \gamma z_1, (w, s)) \text{ for } \gamma \in [0, 1] \quad (5)$$

where γ is a scalar mixing parameter, G is the generator, and (w, s) are word/sentence embedding vectors. In an attempt to quantitatively measure if there is a smooth transition from a perceptually plausible sample to an unsuccessful generated image, we calculate the learned perceptual image patch similarity (LPIPS) [51] score and the perceptual loss [52] which reflect the diversity between two close interpolation samples.

We empirically observe that although the first and last part of interpolation results change gradually with the variations of the latent vectors, both the LPIPS score and the perceptual loss between intermediate samples are the largest and considerably increase, which we detail in Sect. 5.2.1. In other words, when linearly interpolating an unsuccessful latent code and a *Good* latent vector, the appearance and the semantics do not always vary smoothly along with the latent vectors. We therefore make the

assumption that there exists a nonlinear boundary separating successful generated images from unsuccessful synthesized samples in the fake image space. It implicitly means that image quality in the synthetic sample space may be distinguished. Suppose we have a nonlinear image-quality function $f_q : \mathcal{X} \rightarrow t$, where t represents the quality score. We are able to classify a synthesized sample as realistic or unsuccessful.

4.1.2 Good & Bad data set creation

Our goal is to train a powerful classifier that can distinguish successful generated samples from unsuccessful synthetic images. To this end, we built two novel data sets (i.e., the *Good & Bad* bird and face data sets) conditioned on the CUB bird data set [4] and the Multi-Modal CelebA-HQ data set [5], respectively. The *Good & Bad* data set is a collection of perceptually realistic as well as implausible samples generated by a well-trained and fixed text-to-image GAN architecture. The construction of the data set is based on a pilot study [53] on initial manual labeling (210+210 samples), which was used as the training set for automatic good vs bad binary classification. However, such a data set is too small to obtain a training set suitable for end-to-end, deep-learning based quality classification of generated images. Using strict criteria, an extended collection of mixed manual and automatic ‘good’ and ‘bad’ samples was constructed, within one day. Specifically, the used *Good & Bad* bird data set consists of 6700 synthesized samples, i.e., 2700 *Good* and 4000 *Bad* birds. The *Good & Bad* face data set contains 2000 successful generated faces as well as 2000 unsuccessful synthetic faces. A summary of the *Good & Bad* data set is reported in

Table 1 Statistics of the *Good & Bad* bird and face data sets. ‘Bird’ represents the *Good & Bad* bird data set, and ‘Face’ denotes the *Good & Bad* face data set

Data set	Train	Test	Total
Bird	5200	1500	6700
Face	3200	800	4000

Table 1. We visualize a snapshot of our data set in Fig. 4. Our data set is available at https://zenodo.org/record/6283798#.YhkN_ujMI2w. Below, we describe the procedure followed to construct the *Good & Bad* data set.

Image collection The first stage of creating the *Good & Bad* data set involves producing a large set of candidate samples for each data set. The DiverGAN [3] has the ability to adopt a generator/discriminator pair to produce diverse, perceptually plausible and semantically consistent pictures, given a textual description and different injected noise on the latent vector. We therefore choose a pre-trained DiverGAN generator to acquire candidate images. We generated 30,000 synthesized samples as the basis for the selection of a *Good & Bad* bird data set and a *Good & Bad* face data set, respectively.

Image selection Given a variety of candidate pictures, we choose images according to the following criteria:

- (1) A successful generated image is supposed to have vivid shape, rich color distributions, clear background as well as realistic details. For the face data set, photo-realistic images should also have pleasing, undistorted facial attributes (e.g., eyes, hair, makeup, head and mouth) and expressions.
- (2) A synthetic picture with strange shape, blurry background or unclear color is viewed as *Bad*. Meanwhile, we reject faces with an implausible facial appearance or ornamentation (e.g., hat and glasses) as unsuccessful samples.
- (3) We exclude ambiguous images of the type where also for the human judge, the classification as *Good* or *Bad* is difficult. For instance, a bird with only a slightly strange body (e.g., lacking legs) is judged as an ambiguous-quality picture.

For the *Good & Bad* bird data set, we find it inefficient to manually choose thousands of plausible birds from 30,000 collected samples. To reduce the selection labor, we propose a process to obtain the desired birds as follows (depicted in Fig. 5):

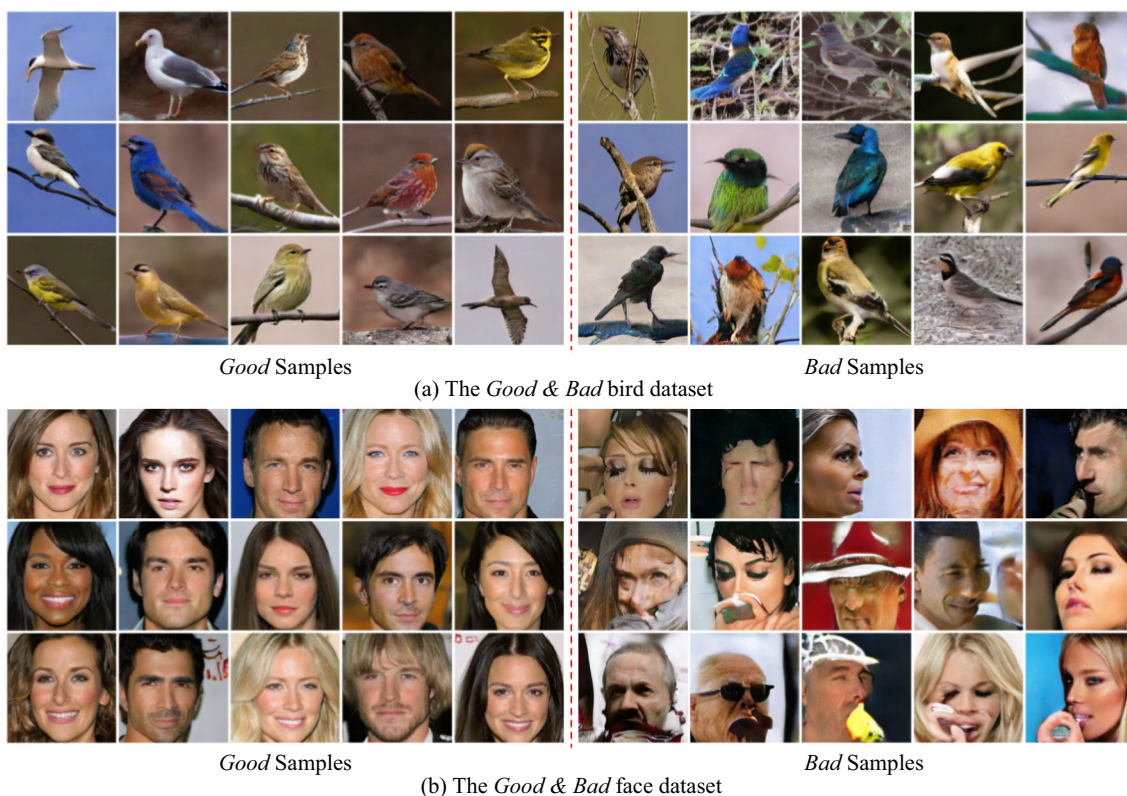


Fig. 4 A snapshot of the *Good & Bad* bird (three top rows) and face (three bottom rows) data sets: The left column is from the *Good* data set; the right column is from the *Bad* data set. These samples are synthesized by the recent DiverGAN generator [3]

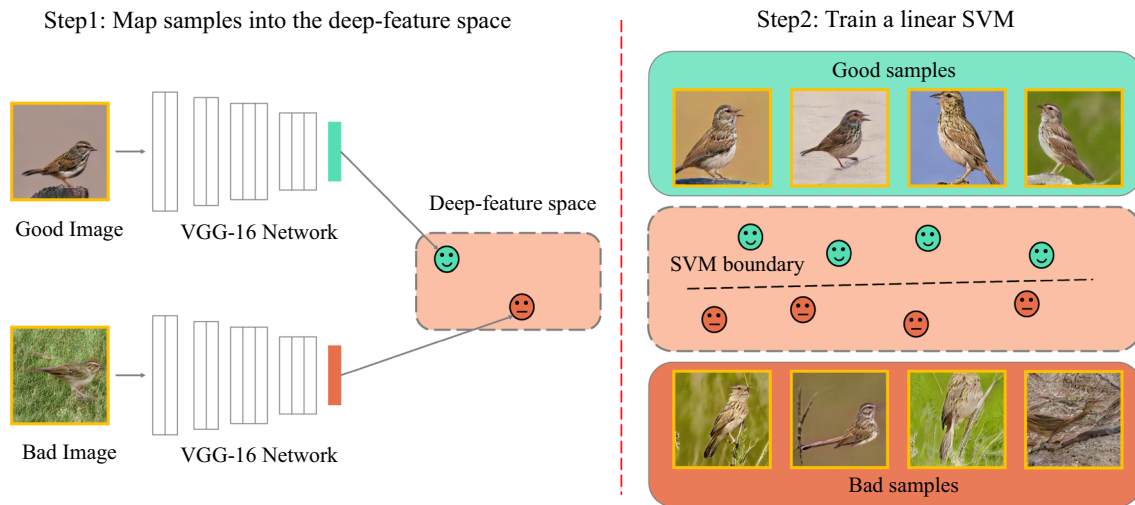


Fig. 5 A schematic outline of the first two steps for automatically discovering *Good* birds from the generated images

- (1) Based on the principles mentioned before, we select 420 synthesized samples (i.e., 210 *Good* and 210 *Bad* birds) as the initial *Good* & *Bad* bird data set, which is split into a training set (i.e., 150 *Good* and 150 *Bad* birds) and a testing set (i.e., 60 *Good* and 60 *Bad* birds). We intend to use these labeled samples to train a simple classification model to try to predict the quality class of synthesized images. However, it is difficult to directly apply a traditional classifier (e.g., a linear SVM) to separate realistic images adequately from inadequate samples, since the image instances exist in a nonlinear manifold [54]. In the meantime, we cannot train a deep neural network (e.g., VGG [55]) from scratch to label a synthetic sample as *Good* or *Bad* due to the small number of the samples in the initial *Good* & *Bad* training set. Bengio et al. [56] postulate that deep convolutional networks have the ability to linearize the manifold of pictures into a Euclidean subspace of deep features. Inspired by this hypothesis, we expect that *Good* and *Bad* samples can be classified by an approximately linear boundary in such deep-feature space.
- (2) We adopt the publicly available VGG-16 network trained on ImageNet to transform the image samples from the training set (i.e., 150 *Good* and 150 *Bad* samples) into the deep-feature representation of layer VGG-16/Conv5_1. The obtained deep features and the corresponding labels (i.e., *Good* and *Bad*) are used to fit a linear SVM model for automatic labeling of the samples in the deep-feature space. To evaluate the performance of the model, we transform the testing samples (i.e., 60 *Good* and 60 *Bad* birds) into deep-feature vectors while applying the learned SVM boundary to predict the classes for the unseen samples.
- (3) In order to harvest an expanded set of *Good* or *Bad* samples, we use the trained SVM model to automatically label the 30,000 collected birds. We manually choose 2700 *Good* and 2000 *Bad* birds from the images that are classified as *Good*, which is not a laborious task due to the performance of the SVM. Moreover, to boost the diversity of *Bad* birds on our data set, we select 2000 *Bad* birds from the samples that are predicted as *Bad*. Finally, 2700 *Good* and 4000 *Bad* birds are acquired as the final, expanded *Good* & *Bad* bird data set. Also for the faces, we discovered that it is easy to label the synthesized samples as *Good* or *Bad*. We therefore manually select 2000 *Good* and 2000 *Bad* samples from 30,000 synthetic faces for the *Good* & *Bad* face data set. The manual selection was realized in one day.

Splitting of the data set The *Good* & *Bad* face data set is randomly divided into the training and test sets with a ratio of 4:1. After the splitting, the training set comprises 3200 images, i.e., 1600 *Good* and 1600 *Bad* faces. The test set consists of 800 samples including 400 *Good* and 400 *Bad* faces. The *Good* & *Bad* bird data set contains 6700 birds, where 5200 images (i.e., 2200 *Good* and 3000 *Bad* birds) belong to the training set and the other 1500 images (i.e., 500 *Good* and 1000 *Bad* birds) belong to the test set.

4.1.3 Synthetic samples classification

Given the extensive training set obtained in this manner, it is now possible to do the quality classification by end-to-end deep learning instead of using an unmodified, pre-trained CNN and an SVM. To fully automatically distinguish successful synthesized samples from unrealistic images, we attempt to fine-tune a pre-trained CNN model

(e.g., ResNet [6]) on the proposed *Good & Bad* data set, which we will detail in Sect. 5.2.3. We expect that this approach is able to achieve the best results. We therefore have the ability to effectively and efficiently identify photo-realistic samples from generated images while acquiring corresponding *Good* latent vectors. These *Good* latent codes can be exploited for further research, facilitating and extending the applicability of text-to-image generation in practice. For instance, we can produce a wealth of high-quality samples by conducting the pairwise linear interpolation between *Good* latent codes, e.g., for the purpose of data augmentation. Given a *Good* latent vector, we can synthesize several similar but semantically diverse pleasing generated samples via latent-space navigation, which will be discussed in the next section.

4.2 Identifying meaningful latent-space directions

In this subsection, we mathematically show that Semantic Factorization (SeFa) [7] approximately identifies the principal components, as PCA does. Furthermore, we propose a technique to capture semantically interpretable latent-space directions for a conditional text-to-image GAN model. To optimize the edited sample, the background-flattening trick is presented to fine-tune the background appearance.

4.2.1 Analyzing the correspondences between SeFa, GANSpace and PCA

We attempt to discuss the relationship between SeFa [7] and GANSpace [8], since they both introduce an algorithmically simple but surprisingly effective technique to derive semantically understandable directions. Specifically, GANSpace collects a set of latent codes and conducts PCA on them to identify the significant latent-space directions. SeFa proposes to directly decompose the pre-trained weights for semantic image editing. Mathematically, SeFa is formulated as in [7]:

$$A^T A n_i - \lambda_i n_i = 0 \quad (6)$$

where $A \in R^{d \times l}$ is the weight matrix of the first transformation step in the generator and $\{n_i\}_{i=1}^k$ indicate k most meaningful directions. The solutions to Eq. 6 correspond to the eigenvectors of $A^T A$ with respect to the k largest eigenvalues. A is usually normalized by L2 norm when implementing SeFa. The formulation of SeFa can almost be perceived as PCA [57] on A , since the results of PCA are the eigen vectors of the covariance matrix C_A associated

with A and C_A is similar to $A^T A$. Specifically, C_A is denoted as in [57]:

$$C_A = \frac{1}{d-1} (A - \langle A \rangle)^T (A - \langle A \rangle) \quad (7)$$

where $\langle A \rangle$ represents the mean from each column of A and C_A is the covariance matrix of A . The difference between regular PCA and SeFa is located in the normalization of A . We therefore argue that SeFa is approximately equivalent to regular PCA on the pre-trained weights. That is to say, GANSpace and SeFa perform PCA on the latent vectors and the pre-trained weights, respectively.

4.2.2 Independent component analysis for semantic discovery in the latent space

It has been observed that the pre-trained weights of the standard GAN contain semantically useful information. We can capture the meaningful latent-space directions in an unsupervised manner by exploiting the well-trained weights of the generator. A conditional text-to-image GAN generator typically leverages a dense layer to transform a latent code into a visual feature map, where a latent space is projected to another space and ultimately into an output image. We make the assumption that there exists a wealth of semantics in the initial fully connected weight matrix of a text-to-image GAN model, due to the linguistic content of the text. We aim at presenting a simple algorithm extracting the main patterns of the pre-trained weights as the interpretable latent-space directions. More specifically, we hypothesize that when given the pre-trained weight matrix A of the first linear layer of the generator, we can obtain a suite of k meaningful semantic factors $N = (n_1, n_2, \dots, n_k)$ by processing the weight matrix A . Mathematically,

$$N = f(A) \quad (8)$$

where $f(\cdot)$ is the function for semantic discovery. These acquired semantics should denote different attributes of the image. For example, n_1 represents pose, n_2 represents smile, and n_3 represents gender for the face data set. To better manipulate the image generation, we argue that these components should be fully independent rather than just uncorrelated (orthogonal). However, when employing PCA as $f(\cdot)$ to discover the controllable latent-space directions, the obtained principal components are only uncorrelated, but not independent. Meanwhile, PCA is optimal for Gaussian data only [10], while the pre-trained weight matrix A is not guaranteed to be Gaussian. Here, we pro-

pose to utilize independent component analysis (ICA) to identify useful latent-space semantics for a conditional text-to-image GAN model.

The goal of ICA is to describe a $M \times L$ data matrix X in terms of independent components. It is denoted as in [10]:

$$X = BS \quad (9)$$

where B is a $M \times T$ mixing matrix and S is a $T \times L$ source matrix consisting of T independent components.

ICA is commonly viewed as a more powerful tool than PCA [58], since it is able to make use of higher-order statistical information incorporating a variety of significant features. Furthermore, ICA is adequate for analyzing non-Gaussian data. To maximize both the independence and the orthogonality between the directions, i.e., n_1, n_2, \dots, n_k , we apply a fast ICA under an additional orthogonality constraint [10] to directly decompose the pre-trained weight matrix to derive the meaningful directions in the latent space. The obtained vectors are therefore not only independent but also orthogonal. We expect that the components can lead to a more precise control over the latent space of the DiverGAN [3] model.

4.2.3 Background flattening

A movement along an effective direction in the latent space should not only accurately change the desired attribute, but also maintain other image content, e.g., the background. However, when applying existing semantic-discovery methods even our introduced algorithm on the text-to-image generation model, we find that the background appearance in the edited sample usually varies along with the target attribute. To overcome this issue, we develop a *Background-Flattening Loss* (BFL) to fine-tune the acquired directions to improve the background. This loss is defined by using both low-level pixels and high-level features, ensuring that the background is optimized and other image contents are preserved. Specifically, it is denoted as:

$$\mathcal{L}_{\text{flatten}}(x_1, x_2) = \|x_1 - x_2\|_1 + \mathcal{L}_{LPIPS}(x_1, x_2) \quad (10)$$

where x_1, x_2 refer to a source sample and an edited sample, respectively. We leverage the Adam algorithm to optimize the independent components.

We empirically find that we are able to employ our proposed BFL to remove the patterns representing the background. To be specific, we can obtain a sample with a white background by increasing the distance (i.e., the BFL) between samples generated by different directions, since the white background and the black background will lead to the maximum loss values. After that, to remove the background, we take the white-background sample as the source image while reducing the distance between the source sample and the edited samples.

4.3 Improving the explainability of the conditional text-to-image GAN

In addition to the latent space, a conditional text-to-image GAN model also contains the linguistic embeddings, in which word and sentence vectors are adopted to modulate the visual feature map for semantic consistency. Despite high-quality pictures achieved by the existing approaches, we yet do not understand what a text-to-image generation architecture has learned within the linguistic space of the conditional input-text probes.

In order to understand ‘embeddings’ in deep learning, several methods have been proposed. A common method is to visualize the space using, e.g., t-SNE or k-means clustering. This may give some insights on the location of dominant image categories in the subspace. An alternative approach is to utilize—yet another—step of dimensionality reduction by applying standard PCA on the embedding. However, this still does not lead to good explanations and an easy controllability of the image-generation process. In this subsection, we start from *Good* latent vectors and introduce two basic techniques to provide insights into the explainability of a text-to-image synthesis framework.

4.3.1 Linear interpolation and semantic interpretability

We study the linear interpolation between a pair of keywords in order to qualitatively explore how well the generator exploits the linguistic space of the conditional input-text probes as well as testing the influence of individual, different words on the generated sample. We can observe how the samples vary as a word in the given text is replaced with another word, for instance by using a polarity axis of qualifier key words (dark-light, red-blue, etc.). More specifically, we can first acquire two word embeddings (i.e., w_0 and w_1) and two corresponding sentence vectors (i.e., s_0 and s_1) by only altering a significant word (e.g., the color attribute value and the background value) in the input natural-language description. Afterward, the results are obtained by performing the linear interpolation between the initial textual description (w_0, s_0) and the changed description (w_1, s_1) while keeping the *Good* latent code z frozen. Mathematically, this proposed text-space linear interpolation combines the latent code, the word and the sentence embeddings and is formulated as:

$$h(\gamma) = G(z, (1 - \gamma)w_0 + \gamma w_1, (1 - \gamma)s_0 + \gamma s_1) \quad (11)$$

where $\gamma \in [0, 1]$ is a scalar mixing parameter and z is a successful latent code.

For the CUB bird data set, when we vary the color attribute value in the given sentence, we empirically explore what happens in the color mix: Do we, e.g., get an

average color interpolation in RGB space or does the network find another solution for the intermediate points between two disparate embeddings?

In general, our presented text-space linear interpolation has the following advantages:

- The linear interpolation between a pair of keywords can be utilized to quantitatively control the attribute of the synthetic sample, when the attribute varies smoothly with the variations of the word vectors. For example, the length of the beak of a bird can be adjusted precisely via the text-space linear interpolation between the word embeddings of ‘short’ and ‘long.’
- When the attribute of the synthesized sample does not change gradually along with the word embeddings, we can exploit a text-space linear interpolation to produce a variety of novel samples. Take bird synthesis as an example: When conducting the linear interpolation between color keywords, the generator is likely to generate a new bird whose body contains two colors (e.g., red patches and blue patches) in the middle of the interpolation results.
- Through the linear interpolation between contrastive keywords, we can take a deep look into which keywords play important roles in yielding foreground images as well as which image (background) regions are determined by the terms in the text probe.

4.3.2 Triangular interpolation and semantic interpretability

We extend the pairwise linear interpolation between two points to the interpolation between three points, i.e., in the 2-simplex, for further studying the generator and better performing data augmentation. Since this kind of interpolation forms a triangular plane, we name it the triangular interpolation. The triangular interpolation is able to generate more and more diverse samples conditioned on three corners (e.g., latent vectors and keywords), spanning a field rather than a line.

Similar to the linear interpolation between a pair of keywords, we need to derive three word embeddings (i.e., w_0 , w_1 and w_2) and three corresponding sentence vectors (i.e., s_0 , s_1 and s_2) as corners to define the presented text-space triangular interpolation:

$$h(\gamma_1, \gamma_2) = G(z, (1 - \gamma_1 - \gamma_2)w_0 + \gamma_1 w_1 + \gamma_2 w_2, (1 - \gamma_1 - \gamma_2)s_0 + \gamma_1 s_1 + \gamma_2 s_2), \quad (12)$$

where $\gamma_1 \in [0, 1]$ and $\gamma_2 \in [0, 1]$ are mixing scalar parameters and z is a successful latent vector.

For the sake of attribute analysis, we can obtain three new textual descriptions by replacing the attribute word in the initial natural-language description with another two

attribute words. Then, through the triangular interpolation between keywords, the generator has the ability to yield pictures based on the above three attributes. Moreover, we expect that the text-space triangular interpolation should achieve the same visual smoothness as the text-space linear interpolation. In other words, when fixing the weight (i.e., γ_2) of the third text in the triangular interpolation between keywords, the attributes of the image vary gradually along with the word embeddings if the interpolation results of a text-space linear interpolation between the first two textual descriptions change continuously.

The text-space triangular interpolation has obvious advantages over the linear interpolation between a pair of keywords. Firstly, the text-space triangular interpolation is able to produce more image variation to perform data augmentation than the pairwise linear interpolation. Secondly, we can simultaneously control two different attributes (e.g., color and the length of the beak) via the triangular interpolation between keywords. Thirdly, through the text-space triangular interpolation, three identical attributes (e.g., red, yellow and blue) can be combined to synthesize a novel sample.

5 Experiments

5.1 Experimental settings

Data sets We perform a set of experiments on three broadly utilized text-to-image data sets, i.e., the CUB bird [4], MS COCO [11] and Multi-Modal CelebA-HQ [5] data sets.

- *CUB bird* The CUB bird data set contains a total of 11,788 images, in which 8855 images are taken as the training set and the remaining 2933 images are employed for testing. Each bird is associated with 10 textual descriptions.
- *MS COCO* The MS COCO data set is a more challenging data set consisting of 123,287 images in total, which are split into 82,783 training pictures and 40,504 test pictures. Each image includes 5 human annotated captions.
- *Multi-Modal CelebA-HQ* The Multi-Modal CelebA-HQ data set is composed of 24,000 and 6000 faces for training and testing, respectively. Each face is annotated with 10 sentences.

Implementation details We take the recent DiverGAN generator [3] as the backbone generator, which is pre-trained on the CUB bird, Multi-Modal CelebA-HQ and MS COCO data sets. The image size of the proposed *Good & Bad* data set is set to $256 \times 256 \times 3$. We set the output dimension of the CNN models (e.g., ResNet [6] and VGG [59]) to 2. We adopt the Adam optimizer with a batch size

of 64 [60] to fine-tune the classification network pre-trained on ImageNet. We utilize the learning-rate finder technique [61] to acquire a suitable learning rate. The one cycle learning rate scheduler is leveraged to dynamically alter the learning rate, while the model is training. We set the mix factor to 3 [7] for SeFa and our proposed algorithm. The scalar parameter for GANSpace [8] is set to 20 on the CUB bird data set and 9 on the COCO data set, respectively. These two parameter values were determined empirically and are needed to compensate for intrinsic scaling differences between the compared methods (Fig. 10). Due to the chosen parameter values, the gradual changes over the horizontal axes of the different approaches are visually more comparable in the figure. We employ the Adam optimizer with $\beta = (0.0, 0.9)$ [37] to fine tune the identified directions. We set the learning rate to 0.0001, as in [37]. The steps of a linear interpolation between two embedded vectors are set to 10. We set the steps of γ_1 and γ_2 in a triangular ‘linguistic’ interpolation to 10. This number is mainly chosen for visual inspection. A smaller number would not show the gradual changes, whereas a larger number would both show a redundant picture and would lead to thumbnail images along the horizontal axis that would be too small. Our methods are implemented by PyTorch. We conduct all the experiments on a single NVIDIA Tesla V100 GPU (32 GB memory).

5.2 Results of finding *Good* synthetic samples

5.2.1 Results of the pairwise linear interpolation of latent codes

To better understand the transition process from a successful synthesized sample to an unsuccessful generated image, we visualize the results of the pairwise linear interpolation between a *Good* latent code and a *Bad* latent vector in Fig. 6a. It can be observed that for the first five and the last two pictures, both the background and the visual appearance of footholds vary gradually along with the latent vectors. However, the background, the visual appearance of footholds, the positions, the shapes and even the orientations (7th \rightarrow 8th sample) of the birds do not change continuously from the 6th image to the 8th sample. It suggests that there may exist a nonlinear boundary separating *Good* samples from *Bad* images in the fake data space.

We also show the corresponding LPIPS score and the perceptual loss (presented in Fig. 6b and c) to quantitatively compare the diversity between two close samples. It can be seen that the increase of the 6th point (6th \rightarrow 7th sample) is the largest and the 7th point (7th \rightarrow 8th

sample) obtains the highest score for both the LPIPS and the perceptual loss. Meanwhile, both points are over the red line which is an approximate boundary distinguishing smooth changes from discontinuous variations and determined by our observations. The results of Fig. 6b and c match what observe in Fig. 6a, indicating that the visual appearance of the birds does not always vary smoothly along with the latent codes.

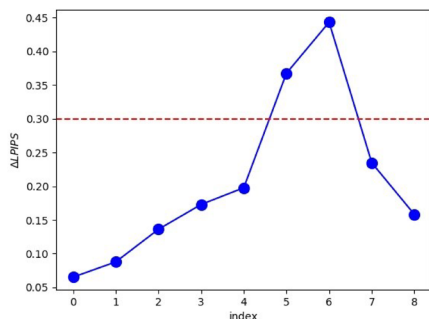
5.2.2 Results on the initial *Good* & *Bad* bird data set

We try different methods to classify a synthetic sample as *Good* or *Bad* on the initial *Good* & *Bad* bird data set (i.e., 210 *Good* and 210 *Bad* birds). The results are reported in Table 2. Here, we discover that all methods using the learned feature vectors of a well-trained VGG-16 network achieve over 94%, suggesting that there exists a (almost) linear boundary in the deep-feature space which can accurately distinguish *Good* samples from *Bad* samples. In addition, the conv5_1 activation in the pre-trained network obtains the best performance (accuracy: 97.5%). We also attempted to employ the SVM with radial basis function (RBF) kernel to classify deep features, acquiring the same result as the linear SVM. Moreover, it can be observed that directly operating on the image pixels (accuracy: 70.0%) and the latent space (accuracy: 75.8%) does not work well for the classification of *Good* and *Bad* samples/latent codes. To boost the accuracy, we conduct PCA on the image pixels to reduce the dimension to 128 and apply a linear SVM to identify realistic samples. However, the accuracy is only improved by 3.3%. The above results confirm the effectiveness of our proposed framework. It should be noted that at 420 test samples, the 95% confidence margin is $\pm 1.5\%$ under the Gaussian or binomial assumption. This means that VGG-16(conv5_3) can be considered as having a worse performance than the two winning variants (Table 2, bottom).

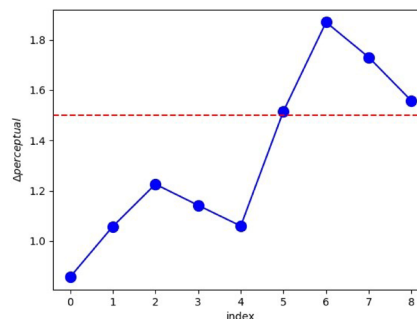
We visualize some typical output samples selected from the test set ($N_{\text{good}}=60$, $N_{\text{bad}}=60$) in Fig. 7 according to their distance to the decision boundary of the trained SVM. It can be observed that *Good* samples are distinguishable from *Bad* samples. Meanwhile, the *Bad* birds around the boundary may have higher quality than the *Bad* birds far from the decision boundary. It should be noted that in non-ergodic problems, where there is not a natural single signal source for the *Good* (or the *Bad*) images, but there rather exists a partitioning of space, the SVM discriminant value for a sample is not guaranteed to be consistent with the intuitive prototypicality of the heterogeneous underlying class [62] due to the lack of a central density for that class.



(a) An example of pairwise linear interpolation of latent codes



(b) The results of the LPIPS score



(c) The results of the perceptual loss

Fig. 6 A random example of a pairwise linear interpolation between latent vectors (left=*Good* → right=*Bad*). The red bounding box in **a** emphasizes a discontinuous range within the linear-interpolation results. The dashed red line in **b** and **c** is an approximate boundary distinguishing smooth changes from discontinuous variations,

determined by our observations. The index number represents the comparison, starting with 0, i.e., the comparison between the first and the second image on the left. The discontinuity is quantitatively revealed both in LPIPS and in perceptual loss (Color figure online)

Table 2 Classification accuracy on the separation boundary with respect to image quality.

Methods	Accuracy (%)
Image	70.0
PCA-Image	73.3
Latent Code	75.8
VGG-16(conv5_3)	94.2
VGG – 16(conv5_2)	96.7
VGG – 16(conv5_1)	97.5

PCA-Image refers to using PCA on the image pixels after reducing the dimensionality to 128 and applying SVM to identify realistic samples. *Latent Code* refers to the direct application of SVM in the latent space. *Image* refers to a direct application of SVM on the image pixels. At 420 test samples, the 95% confidence margin is ±1.5%, allowing to indicate the best two methods (bold)

5.2.3 Results on the *Good* & *Bad* data set

The classification results We fine-tune the pre-trained CNN models (i.e., ResNet and VGG) on the *Good* & *Bad* data set in order to accurately predict the quality classes of generated images. The comparison between VGG-11, VGG-16, VGG-19, ResNet-18, ResNet-50 and ResNet-101 with respect to the classification performance on the *Good* & *Bad* bird and face data sets is shown in Table 3. We can observe that ResNet-50 achieves the best result (accuracy: 98.09%) on the *Good* & *Bad* bird data set and ResNet-101 impressively acquires the accuracy of 99.16% on the *Good* & *Bad* face data set. It can also be seen that ResNet performs better than VGG and all the networks

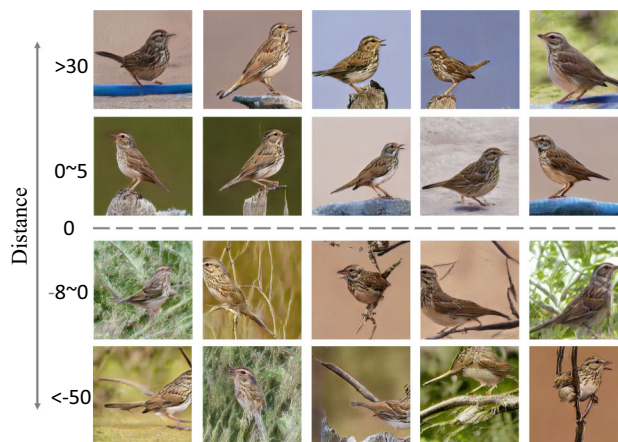


Fig. 7 Example of partitioning of latent-code space between *Good* (two top rows) and *Bad* latent codes (two bottom rows), as determined by the discriminant value (distance) computed by a linear SVM (training set: $N_{good}=150$, $N_{bad}=150$.)

obtain a better than 95% accuracy on both the *Good* & *Bad* bird data set and the *Good* & *Bad* face data set. The above results demonstrate that the *Good* and *Bad* samples in the synthetic image space can be effectively distinguished by a well-trained deep convolutional network. It should be noted that at 4700 test samples, the 95% confidence margin is ±0.45% under the Gaussian or binomial assumption. This implies that for faces (Table 3, right column), Resnet101 is the clear winner, whereas for birds (Table 3, left column), VGG-19, ResNet-18, ResNet-50 and Res-Net-101 have a similar performance, statistically.

Table 3 Classification performance of the deep convolutional networks on the *Good & Bad* bird and face data sets. *Bird* refers to the *Good & Bad* bird data set, and *Face* refers to the *Good & Bad* face data set

Method	Classification performance/%	
	Data set	
	Bird	Face
VGG-11	96.53	98.08
VGG-16	95.70	97.84
VGG-19	97.85	98.20
ResNet-18	97.59	98.68
ResNet-50	98.09	98.56
ResNet-101	97.79	99.16

The best results are in bold

Visualization of the learned representation To visually investigate the distribution of the features learned by the CNN models (i.e., ResNet-50 for the *Good & Bad* bird data set and ResNet-101 for the *Good & Bad* face data set), we exploit the PCA and t-SNE [64] approaches to embed the samples on the *Good & Bad* data set into a 2-dimensional space as shown in Fig. 8. From the scatter plots in this figure, we can see that the learned representations of the classification networks from different classes (i.e., *Good* and *Bad*) are well separated, indicating that the image classification models can project the plausible vs unrealistic samples into two different regions of the latent spaces. Therefore, discovering photo-realistic samples from synthesized images appears to be feasible. It can also be observed that *Good & Bad* samples in the **faces** data set are more separable in PCA than the *Good & Bad* samples in the **bird** data set, which may explain why ResNet-101 trained on the face data set performs better than ResNet-50 trained on finding good vs bad samples in the bird data set. The latter is consistent with the obtained classification score.

Explaining the classification prediction We leverage three different methods (i.e., Layer-CAM [65], integrated gradient [66] and extremal perturbation [67]) to explain the image classification prediction obtained by ResNet-50

trained on the *Good & Bad* bird data set and ResNet-101 trained on the *Good & Bad* face data set. Figure 9 shows the explanation for the top 1 predicted class, suggesting that the classification network derives the results by concentrates on the discriminative regions of the objects (i.e., birds and face). For instance, Layer-CAM visualization (2nd and 6th column) localizes the heads and belly of the birds and the noses, mouths and eyes of the faces. Meanwhile, integrated gradient (3rd and 7th column) and extremal perturbation (4th and 8th column) correctly highlight the branches and the whole bodies of the birds while capturing the hat and the entire faces, pinpointing the reason why the samples are classified into the corresponding categories. More importantly, the blurry regions of the images (6th, 7th and 8th column) are accurately identified by these explainable approaches. That is to say, our classification model can separate implausible regions from high-quality patches and discover successful synthetic samples from generated images.

Quantitative comparison of image quality We evaluate the image quality of the proposed *Good & Bad* data set quantitatively. We randomly select 2000 *Good* images and 2000 *Bad* images from the data set to compute the Fréchet inception distance (FID) score, which is reported in Table 4. We can see that *Good* samples achieve significantly lower FID scores than *Bad* samples on the *Good & Bad* bird and face data sets, which suggests the usefulness of the proposed method. Since the inception score (IS) may be saturated, even over-fitted [3, 36, 37] and we do not have enough pictures, we do not compare the IS on our data set.

5.3 Results of latent-space manipulation

5.3.1 Comparison between SeFa, GANSpace and PCA

Figure 10 plots the latent-code manipulation results of SeFa [7], GANSpace [8] and regular PCA on the CUB bird

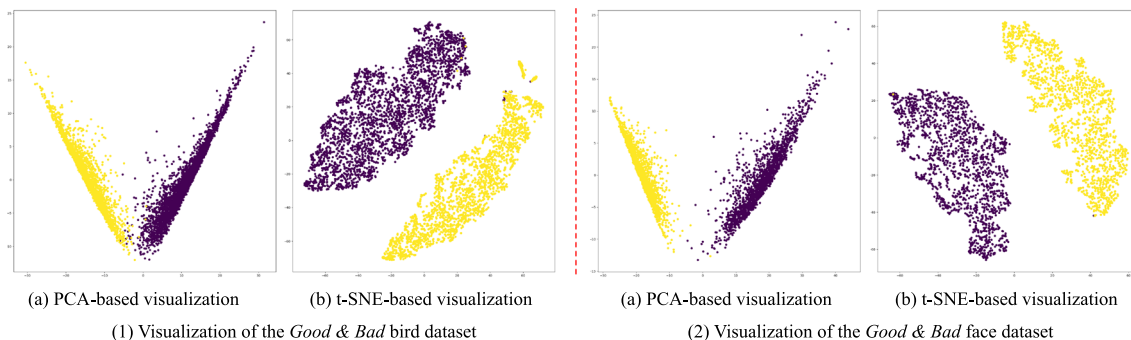


Fig. 8 Separability of the *Good & Bad* classes, left **a** for the birds, right **b** for the face-image data set, by utilizing the PCA and t-SNE [63] methods. These scatter plots show the good separability of the

two qualitative classes: The yellow color represents the *Good* sample, while purple dots represent *Bad* image samples (Color figure online)

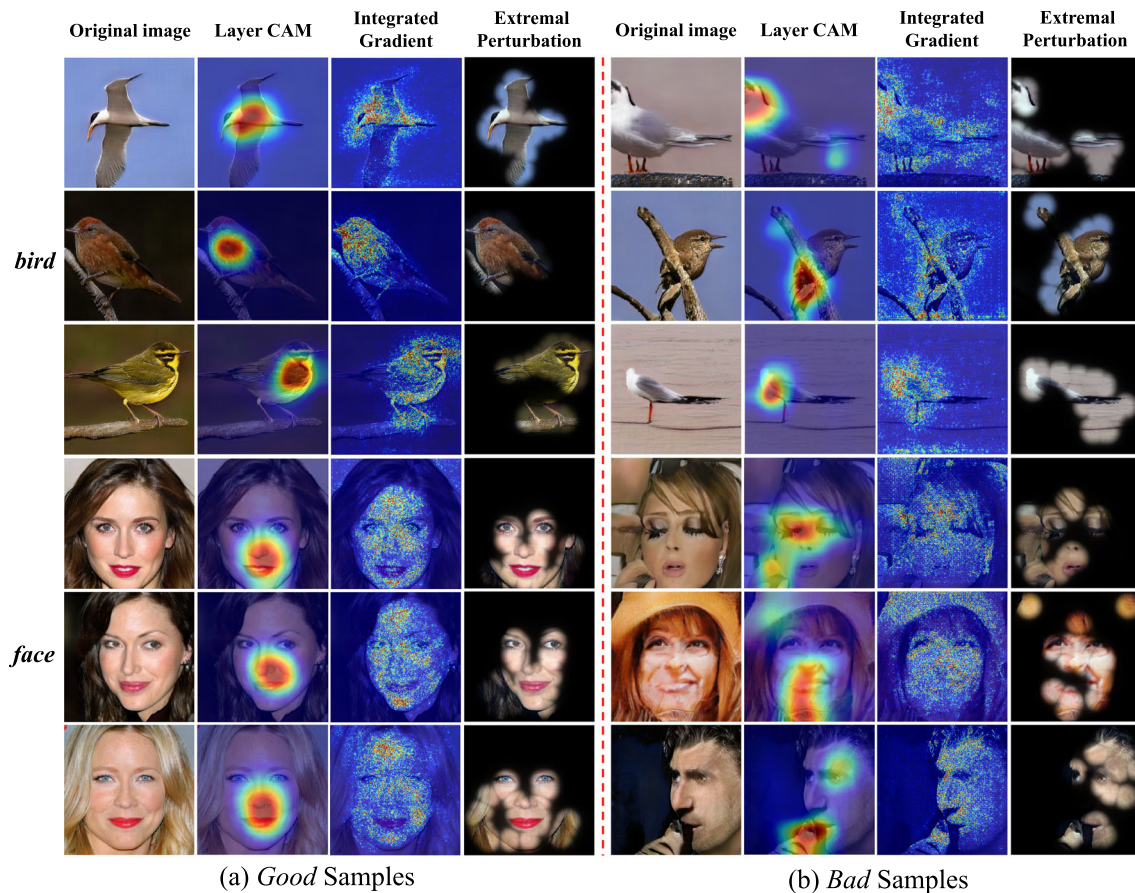


Fig. 9 Explaining the image classification prediction made by ResNet-50 on the *Good* & *Bad* bird data set (three top rows) and ResNet-101 on the *Good* & *Bad* face data set (three bottom rows) using Layer-CAM [65], integrated gradient [66] and extremal

perturbation [67]. The **left** half of the grid is from the *Good* data set; the **right** half of the grid is from the *Bad* data set, separated by the dashed line

Table 4 The FID of *Good* pictures and *Bad* pictures on the *Good* & *Bad* data set. A lower FID score means that the generated pictures are closer to the corresponding real pictures. *Bird* refers to the *Good* & *Bad* bird data set, and *Face* refers to the *Good* & *Bad* face data set

Data set	Good	Bad
Bird	38.28	47.39
Face	30.93	45.62

The best results are in bold

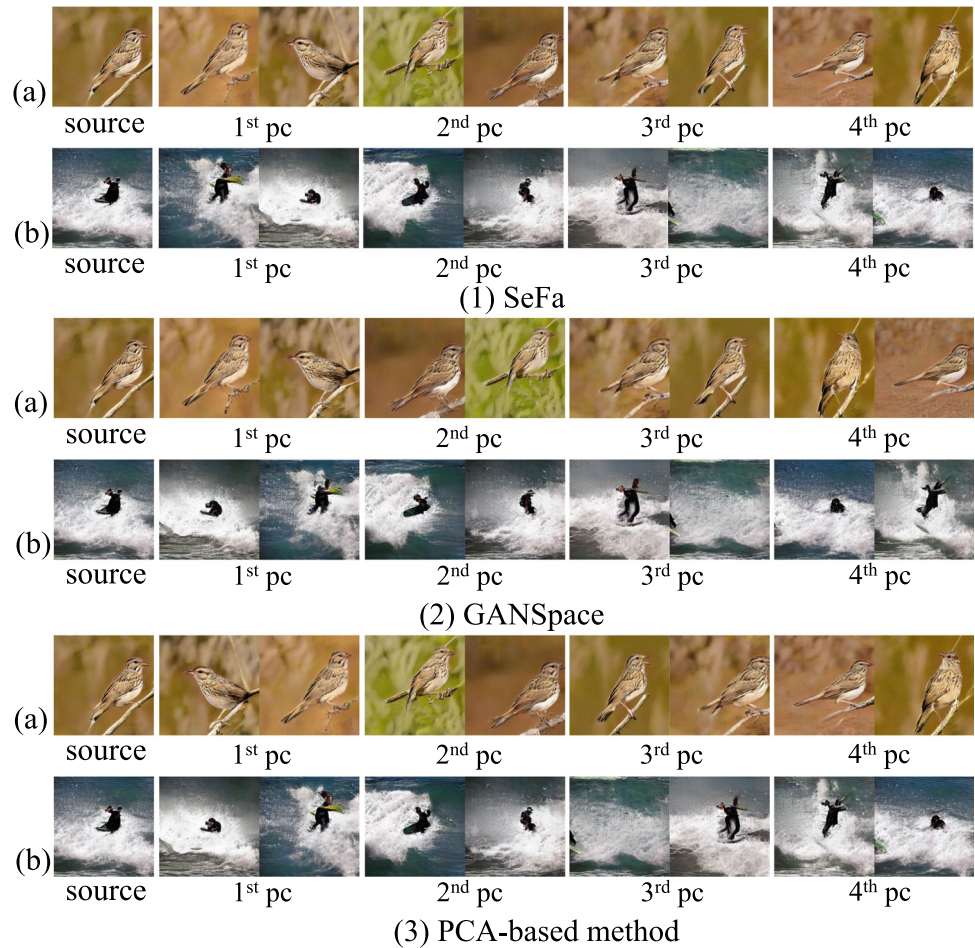
and COCO data sets. We discover that these three approaches derive almost the identical directions although for some components (e.g., 4th principal component) the negative and the positive side is reversed, supporting our claim in Sect. 4.2.1. Note that GANSpace is implemented by leveraging the first dense layer of DiverGAN to collect 10,000 sets of feature maps while performing PCA on them to obtain principal components as useful attributes. Additionally, we adjust the max mix factor to 20 on the CUB bird data set and 9 on the COCO data set, respectively. The

above analysis suggests that when enough data are sampled, SeFa is similar to GANSpace for DiverGAN.

5.3.2 Comparison with unsupervised methods

For qualitative comparison, we visualize the meaningful directions identified by our proposed algorithm and SeFa on the CUB bird and Multi-Modal CelebA-HQ data sets in Fig. 11. We can tell that our method is able to derive several fine-grained semantics corresponding to rotation, background and size for the bird model and pose, hair and smile for the face model, validating its effectiveness. Meanwhile, our approach leads to a more powerful control over the latent codes than SeFa. For example, when editing the size of the bird, our algorithm better preserves the background appearance. Our method also achieves better performance than SeFa from the perspective of smiling. It can also be seen that our method captures the same rotation and pose attributes as SeFa. The reason for this may be that ICA under orthogonal constraint and PCA can discover

Fig. 10 Visualization of individual components within the latent codes, for (1) SeFa [7], (2) GANSpace [8] and (3) regular PCA. The original source image is in the left column (2 examples, a and b). For each principal component (pc1-pc4), example images from the negative and the positive side of its axis are shown



exactly the same most representative semantics (rotation for the bird model and pose for the face model).

We discovered that our algorithm is able to better remove the smile than SeFa, which is illustrated in Fig. 12. We can observe that our method can edit the smile on the face while preserving other attributes, such as pose, background and hair. However, when removing the smile using SeFa, background and hair style are changed.

We perform quantitative analysis on two captured semantics, i.e., smile and size, in order to further validate the effectiveness of our algorithm. To quantitatively evaluate the direction about smile identified by our approach, following [9], we train a smiling classifier on the Multi-Modal CelebA-HQ data set with ResNet-50 network [6]. The result of the classifier may indicate whether the discovered direction is correlated with the desired attribute. Note that since many significant attributes, such as gender, age and glasses, are only manipulated by textual descriptions, we do not train the corresponding predictors. For the quantitative analysis of the direction about the size of the bird, inspired by [45], we use the cosine distance between

the CLIP embeddings of the edited birds and the text ‘a large bird’ as the evaluation metric. The reason behind using the CLIP loss is that it can effectively guide image manipulation through a text prompt. The results are depicted in Table 5. We can see that our algorithm outperforms SeFa on both the smile and the size of the bird, which suggests its effectiveness. The above results demonstrate that based on the *Good* latent codes found by our well-trained classification model, we can adopt our presented algorithm to acquire a wealth of semantically diverse and perceptually realistic samples.

5.3.3 Human evaluation

We conduct a human-preference test on the Multi-Modal CelebA-HQ data set to compare our method with SeFa. We randomly select 100 successful synthesized faces while employing the directions (i.e., smile and hair) found by these two approaches. The user is asked to control the relevant dimension using a slider, for both methods. Users

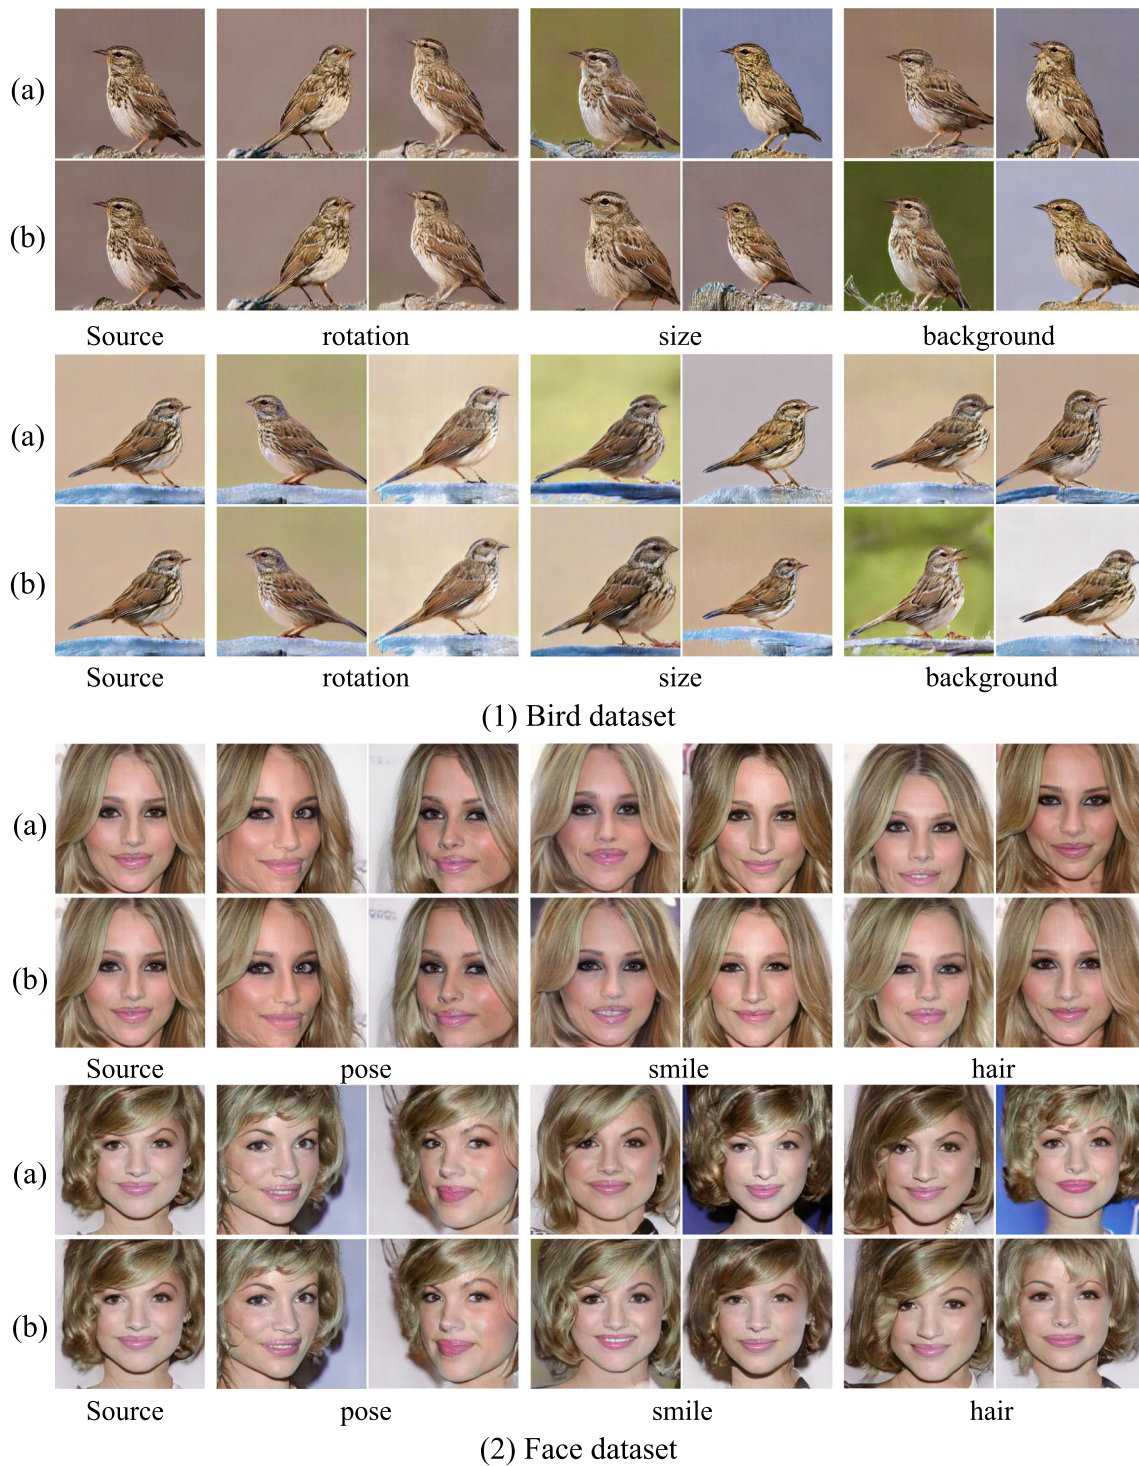


Fig. 11 Qualitative comparison of the meaningful latent-space directions discovered by **a** SeFa [7] and **b** our proposed algorithm on (1) the CUB bird (four top rows) and (2) Multi-Modal CelebA-HQ (four bottom rows) data sets

are then asked to choose the sample with the most accurate change. Two human subjects participated in this evaluation. The performance indicator is the percentage of cases where a method was selected as the best. As illustrated in

Fig. 13, our method performs better than SeFa with respect to the control of smile and hair, which demonstrates the superiority of our proposed algorithm.

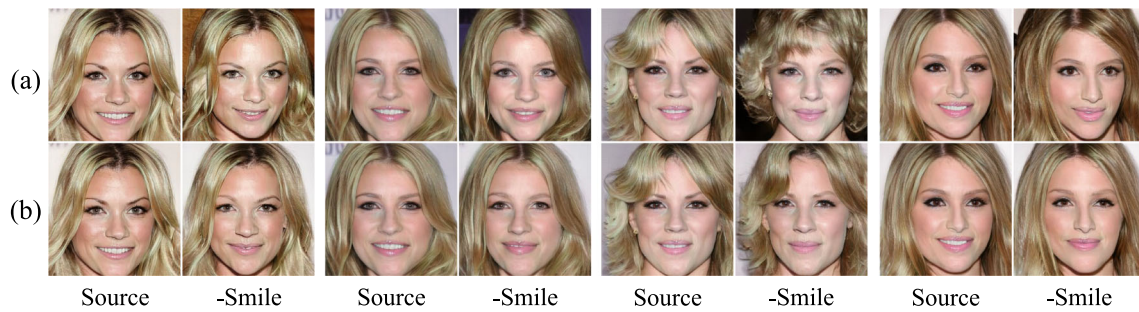


Fig. 12 Qualitative comparison of the removing-smile latent-space direction discovered by **a** SeFa [7] and **b** our proposed algorithm on the Multi-Modal CelebA-HQ data set. *–Smile* indicates the removing-smile image

Table 5 Quantitative analysis of the size and smiling semantics discovered by our method and SeFa [7]

Method	Size ↑	Smile ↑
SeFa	0.25	0.54
Ours	0.27	0.56

The best scores are in bold

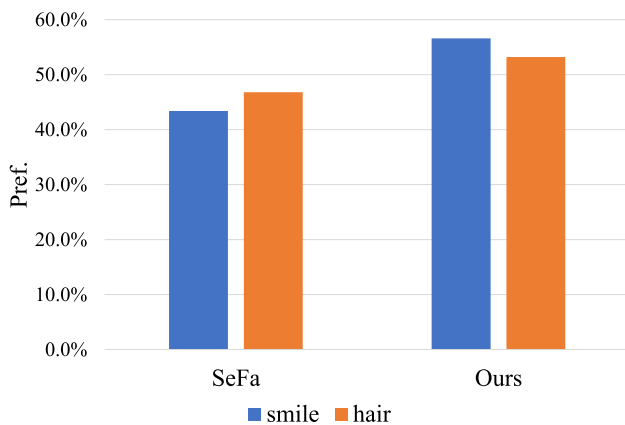


Fig. 13 Human-rater preference (Pref.) for the method SeFa [7] and our proposed method with respect to the task ‘control the smile’ (blue bars) and the task ‘control the hair’ (orange bars) on the Multi-Modal CelebA-HQ data set. The y axis represents the percentage of trials where a method was chosen. The sum of the voting percentages over methods is 100% within a task. The total number of trials per task is 100, i.e., 200 votes were obtained (Color figure online)

5.3.4 Results of background flattening

To prove the effectiveness of background flatten, we apply it to optimize the directions obtained by our proposed algorithm. The results are illustrated in Fig. 14. By comparing the first row with the second row, we can see that the background is significantly improved and other image contents are maintained, indicating that the presented background-flattening method can be employed for existing latent-code manipulation approaches to fine-tune the backgrounds of synthetic samples. As can be observed in the third row, background flattening can also be leveraged

to remove the background while keeping the birds unchanged.

5.4 Results of a ‘linguistic’ interpolation

5.4.1 Results of the linear interpolation between keywords

Figure 15 shows the qualitative results of the linear ‘linguistic’ interpolation of DiverGAN on the CUB bird data set, indicating that the attributes correlated with the synthesized sample do not always change gradually with the variations of word embeddings. For instance, the color of the bird does not vary continuously from ‘red’ to ‘blue’ in the first row. In the medium of interpolation results, DiverGAN generates multiple novel birds, whose bodies are composed of red and blue patches. However, the color attribute of the bird changes gradually from ‘red’ to ‘yellow’ in the second row. We are able to acquire an average color interpolation in RGB space by merging the first and second attributes. We can also see that in the third row, the length of the beak varies smoothly along with textual vectors, while other attributes remain unchanged. Furthermore, while the color of the beak changes continuously with the variations of word embeddings, the shape of the bird varies largely in the fourth row. The above results suggest that DiverGAN has the ability to capture the significant words (e.g., the color of the body and the length of the beak) in the given textual description. More importantly, by exploiting the characteristic as well as the linear interpolation between a pair of keywords, we can precisely control the image-generation process while producing various novel samples.

The qualitative results of the linear interpolation between contrastive keywords on the COCO data set are shown in Fig. 16. We can observe that DiverGAN accurately identifies ‘beach,’ ‘snow’ and ‘men’ while generating the corresponding image samples. In addition, the background (1st and 2nd row) and the object (3rd row) change continuously along with linguistic vectors. It can also be seen that although we change the ‘acting’ word

Fig. 14 Visualization of background flatten and background removal for the meaningful directions acquired by our proposed method

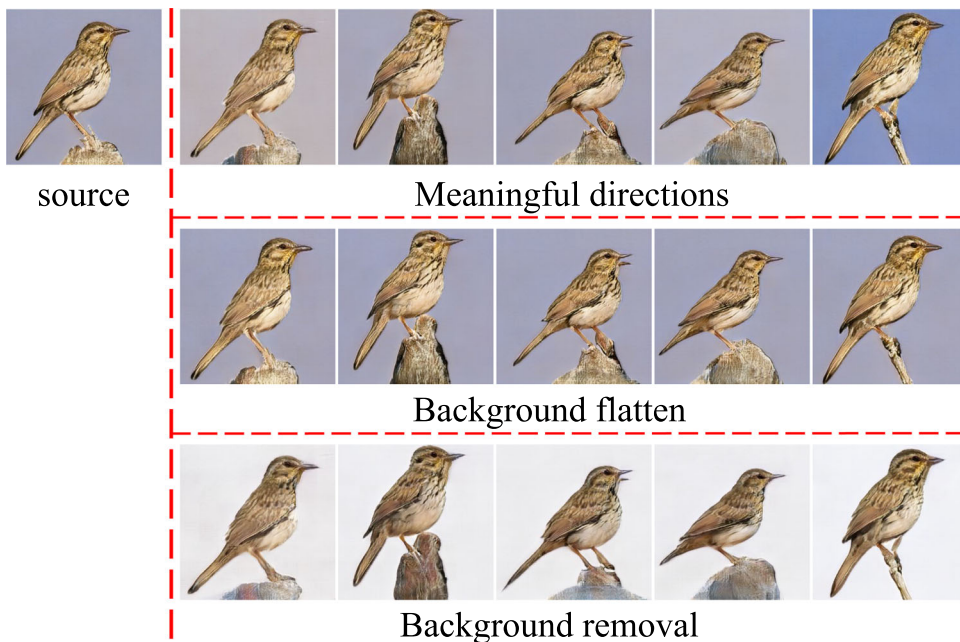


Fig. 15 ‘Linguistic’ interpolation of DiverGAN random latent-code samples on the CUB data set, for four text input probes

from ‘grazing’ to ‘skiing,’ the background significantly varies from ‘grass’ to ‘snow’ in the fourth row, which demonstrates that some words (e.g., ‘skiing’) play a vital role in the generation process of image samples. Furthermore, the above analysis indicates that when given adequate training images, DiverGAN is able to control the background (e.g., from grass to beach) and object (e.g., from animals to men) of complex scenes with the help of the linear ‘linguistic’ interpolation, since DiverGAN is able to learn the corresponding semantics in the linguistic space of the conditional input-text probes.

In addition to visualizing effective examples of the linear interpolation between keywords, we also present

some unsuccessful results in Fig. 17. As can be observed in Fig. 17, the size of the bird (1st and 2nd row) does not vary with the variations of the word (from ‘small’ to ‘big’ and from ‘small’ to ‘medium’). In addition, we can see that the background (3rd row) and the object (4th row) unfortunately do not change along with the word (from ‘grass’ to ‘street’ and from ‘animals’ to ‘cows’). At this point we can conclude that many meaningful contrasts can be learned (Fig. 16), but there are areas where the method is not able to capture important variations along a dimension. This may be due to architectural or data-related limitations. In order to improve our insights, we will look at a triangular interpolation in the next subsection.

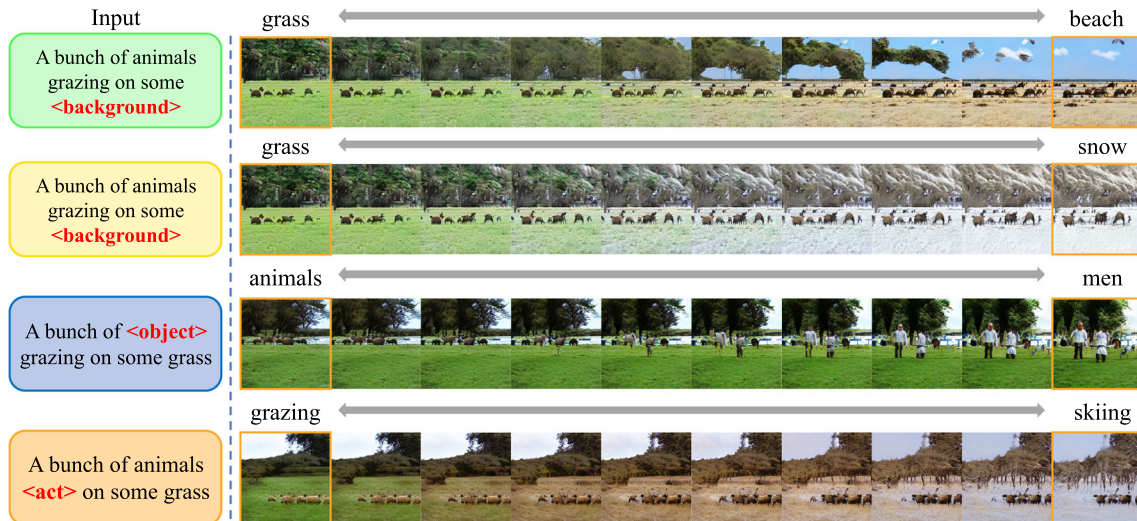


Fig. 16 ‘Linguistic’ interpolation of DiverGAN random latent-code samples on the COCO data set, for four text input probes



Fig. 17 Unsuccessful ‘linguistic’ interpolation of DiverGAN random latent-code samples on the CUB and COCO data sets, for four text input probes. For the third row, the desired attribute (i.e., a street) is not emerging

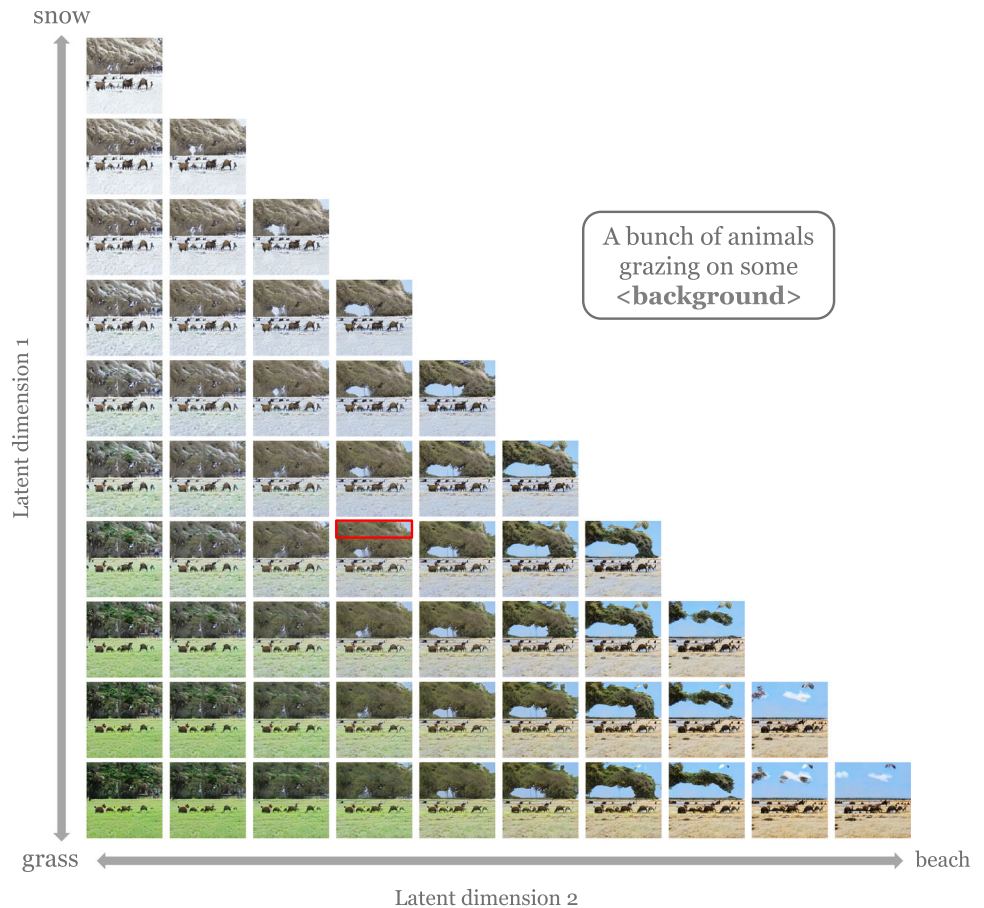
5.4.2 Results of a triangular ‘linguistic’ interpolation

The triangular interpolation for linguistic attributes (i.e., the points between *snow*, *grass*, *beach*) in two dimensions is shown in Fig. 18. We can observe that the transitions toward the three corner points are natural as well as smooth. Furthermore, the interpolation results achieve a balanced triangular shape within the triangle, such that the center marked in red is the combination of three linguistic attributes. If the application concerns data augmentation, 55 believable samples are obtained by performing the triangular interpolation between keywords.

6 Conclusion

In this paper, we propose several techniques to overcome the challenges of text-to-image generation in real-world applications. To ensure the quality of synthetic pictures, we created a *Good & Bad* data set, both for a bird and a face-image collection, which comprises high-resolution as well as implausible synthesized samples, in which the images are chosen by following strict principles. Based on the *Good & Bad* data set, we fine-tune the deep convolutional network trained on ImageNet to classify a generated image as *Good* or *Bad*. To better understand and exploit the latent space of a conditional text-to-image GAN model, we introduce the independent component analysis (ICA) algorithm under an additional orthogonal constraint that

Fig. 18 The triangular interpolation of latent codes, for linguistic attributes *snow*, *grass*, *beach* on two dimensions. The center is marked in red (Color figure online)



can extract both independent and orthogonal components from the pre-trained weight matrix of the generator as the semantically interpretable latent-space directions. In addition, we designed a background-flattening loss (BFL) to optimize the background appearance in the edited sample. To provide valuable insight into the relationship between the linguistic embeddings and the synthetic-sample semantic space, we conduct linear interpolation analysis between pairs of keywords. Meanwhile, we extend a pairwise linear interpolation to a triangular interpolation conditioned on three corners to further analyze the model.

We evaluate our presented approaches on the recent DiverGAN generator that was pre-trained on three popular data sets, i.e., the CUB bird, Multi-Modal CelebA-HQ and MS COCO data sets. Extensive experimental results suggest that our well-trained classifier is able to accurately predict the quality classes of the samples from the testing set and our introduced algorithm can derive meaningful semantic properties in the latent space of DiverGAN, which validates the effectiveness of our proposed methods. Furthermore, we show that semantics contained in the image change gradually with the variations of latent codes, but the attributes of the sample do not always vary continuously along with the word embeddings: For some

contrasts, there may exist fast transitions along the latent axis. This may not be surprising because not all semantic categories are guaranteed to have a fluent transition. In many applications, size and relative size of objects will be considered as an important control variable. However, we find that DiverGAN, being a convolutional neural network, cannot capture the size of objects. Notably, this invariance is desirable in a classification context, but not optimal in a generative application. Furthermore, DiverGAN cannot understand some words in the given textual description. This may be due to the lack of (diverse) textual descriptions in the labels accompanying the images in the data sets, a common problem in generative text-to-image approaches. The advent of large language models may help in diversifying the textual annotations in image collections. In the future, we will explore how to utilize the presented approach to perform data augmentation for training image classifiers. Instead of a Monte Carlo like random sampling of augmentation candidate images, our method allows for a more targeted sampling procedures, taking into account semantic contrasts and/or the estimated quality of generated images.

Acknowledgements We would like to thank the Center for Information Technology of the University of Groningen for their support and

for providing access to the Peregrine high-performance computing cluster.

Data Availability The data sets generated during and/or analyzed during the current study are available from the corresponding author on reasonable and motivated request.

Declarations

Conflict of interest The authors declared that they have no conflicts of interest in this work.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Mirza M, Osindero S (2014) Conditional generative adversarial nets. arXiv preprint [arXiv:1411.1784](https://arxiv.org/abs/1411.1784)
- Zhang Z, Schomaker L (2022) Optgan: Optimizing and interpreting the latent space of the conditional text-to-image gans. arXiv preprint [arXiv:2202.12929](https://arxiv.org/abs/2202.12929)
- Zhang Z, Schomaker L (2021) Divergan: an efficient and effective single-stage framework for diverse text-to-image generation. *Neurocomputing* 473:182–198
- Wah C, Branson S, Welinder P, Perona P, Belongie S (2011) The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology
- Xia W, Yang Y, Xue J-H, Wu B (2021) Tedigan: text-guided diverse face image generation and manipulation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2256–2265
- He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778
- Shen Y, Zhou B (2021) Closed-form factorization of latent semantics in gans. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1532–1540
- Härkönen E, Hertzmann A, Lehtinen J, Paris S (2020) Ganspace: Discovering interpretable gan controls. arXiv preprint [arXiv:2004.02546](https://arxiv.org/abs/2004.02546)
- Shen Y, Gu J, Tang X, Zhou B (2020) Interpreting the latent space of gans for semantic face editing. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9243–9252
- Ablin P, Cardoso J-F, Gramfort A (2018) Faster ica under orthogonal constraint. In: *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4464–4468. IEEE
- Lin T-Y, Maire M, Belongie S, Hays J, Perona P, Ramanan D, Dollár P, Zitnick CL (2014) Microsoft coco: Common objects in context. In: *European Conference on Computer Vision*, pp. 740–755. Springer
- Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y (2014) Generative adversarial nets. *Adv Neural Inf Process Syst* 27:2672–2680
- Gao L, Chen D, Zhao Z, Shao J, Shen HT (2021) Lightweight dynamic conditional GAN with pyramid attention for text-to-image synthesis. *Pattern Recognit* 110:107384
- Shi H, Wang L, Zheng N, Hua G, Tang W (2022) Loss functions for pose guided person image generation. *Pattern Recognit* 122:108351
- Yan L, Zheng W, Gou C, Wang F-Y (2021) ISGAN: identity-sensitive generative adversarial network for face photo-sketch synthesis. *Pattern Recognit* 119:108077
- Zhang X, Wang X, Shi C, Yan Z, Li X, Kong B, Lyu S, Zhu B, Lv J, Yin Y et al (2021) De-GAN: domain embedded GAN for high quality face image inpainting. *Pattern Recognit* 124:108415
- Yang F, Ren J, Lu Z, Zhang J, Zhang Q (2022) Rain-component-aware capsule-GAN for single image de-raining. *Pattern Recognit* 123:108377
- Goodfellow I, Warde-Farley D, Mirza M, Courville A, Bengio Y (2013) Maxout networks. In: *International Conference on Machine Learning*, pp. 1319–1327. PMLR
- LeCun Y, Bottou L, Bengio Y, Haffner P (1998) Gradient-based learning applied to document recognition. *Proc IEEE* 86(11):2278–2324
- Denton E, Chintala S, Szlam A, Fergus R (2015) Deep generative image models using a Laplacian pyramid of adversarial networks. arXiv preprint [arXiv:1506.05751](https://arxiv.org/abs/1506.05751)
- Radford A, Metz L, Chintala S (2015) Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint [arXiv:1511.06434](https://arxiv.org/abs/1511.06434)
- Ioffe S, Szegedy C (2015) Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *International Conference on Machine Learning*, pp. 448–456. PMLR
- Zhang H, Goodfellow I, Metaxas D, Odena A (2019) Self-attention generative adversarial networks. In: *International Conference on Machine Learning*, pp. 7354–7363. PMLR
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017) Attention is all you need. In: *Advances in Neural Information Processing Systems*, pp. 5998–6008
- Brock A, Donahue J, Simonyan K (2018) Large scale gan training for high fidelity natural image synthesis. arXiv preprint [arXiv:1809.11096](https://arxiv.org/abs/1809.11096)
- Karras T, Laine S, Aila T (2019) A style-based generator architecture for generative adversarial networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4401–4410
- Kumar A, Alsadoon A, Prasad P, Abdullah S, Rashid TA, Pham DTH, Nguyen TQV (2022) Generative adversarial network (GAN) and enhanced root mean square error (ERMSE): deep learning for stock price movement prediction. *Multimed Tools Appl* 81:3995–4013
- Jameel SK, Aydin S, Ghaeb NH, Majidpour J, Rashid TA, Salih SQ, JosephNg PS (2022) Exploiting the generative adversarial network approach to create a synthetic topography corneal image. *Biomolecules* 12(12):1888
- Zhang H, Xu T, Li H, Zhang S, Wang X, Huang X, Metaxas DN (2017) StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5907–5915
- Xu T, Zhang P, Huang Q, Zhang H, Gan Z, Huang X, He X (2018) AttnGAN: Fine-grained text to image generation with

- attentional generative adversarial networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1316–1324
31. Qiao T, Zhang J, Xu D, Tao D (2019) MirrorGAN: Learning text-to-image generation by redescription. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1505–1514
 32. Zhu M, Pan P, Chen W, Yang Y (2019) Dm-GAN: Dynamic memory generative adversarial networks for text-to-image synthesis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 5802–5810
 33. Hinz T, Heinrich S, Wermter S (2019) Semantic object accuracy for generative text-to-image synthesis. arXiv preprint [arXiv:1910.13321](https://arxiv.org/abs/1910.13321)
 34. Cheng J, Wu F, Tian Y, Wang L, Tao D (2020) Rifegan: Rich feature generation for text-to-image synthesis from prior knowledge. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 10911–10920
 35. Radford A, Metz L, Chintala S (2016) Unsupervised representation learning with deep convolutional generative adversarial networks. CoRR [arXiv:1511.06434](https://arxiv.org/abs/1511.06434)
 36. Tao M, Tang H, Wu S, Sebe N, Wu F, Jing X-Y (2020) DF-GAN: Deep fusion generative adversarial networks for text-to-image synthesis. arXiv preprint [arXiv:2008.05865](https://arxiv.org/abs/2008.05865)
 37. Zhang Z, Schomaker L (2021) Dtgan: Dual attention generative adversarial networks for text-to-image generation. In: 2021 International Joint Conference on Neural Networks (IJCNN), pp. 1–8. IEEE
 38. Zhang H, Koh JY, Baldrige J, Lee H, Yang Y (2021) Cross-modal contrastive learning for text-to-image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 833–842
 39. Wang H, Lin G, Hoi SC, Miao C (2021) Cycle-consistent inverse gan for text-to-image synthesis. In: Proceedings of the 29th ACM International Conference on Multimedia, pp. 630–638
 40. Afifi S, Gholamhosseini H, Sinha R (2020) Dynamic hardware system for cascade svm classification of melanoma. *Neural Comput Appl* 32(6):1777–1788
 41. Goetschalckx L, Andonian A, Oliva P (2019) Ganalyze: Toward visual definitions of cognitive image properties. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 5744–5753
 42. Voynov A, Babenko A (2020) Unsupervised discovery of interpretable directions in the gan latent space. In: International Conference on Machine Learning, pp. 9786–9796. PMLR
 43. Jahanian A, Chai L, Isola P (2019) On the steerability of generative adversarial networks. arXiv preprint [arXiv:1907.07171](https://arxiv.org/abs/1907.07171)
 44. Wang H-P, Yu N, Fritz M (2021) Hijack-gan: Unintended-use of pretrained, black-box gans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7872–7881
 45. Patashnik O, Wu Z, Shechtman E, Cohen-Or D, Lischinski D (2021) Styleclip: Text-driven manipulation of stylegan imagery. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 2085–2094
 46. Radford A, Kim JW, Hallacy C, Ramesh A, Goh G, Agarwal S, Sastry G, Askell A, Mishkin P, Clark J, et al (2021) Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning, pp. 8748–8763. PMLR
 47. Schuster M, Paliwal KK (1997) Bidirectional recurrent neural networks. *IEEE Trans Signal Process* 45(11):2673–2681
 48. Sarkar A, Hossain SS, Sarkar R (2023) Correction to: human activity recognition from sensor data using spatial attention-aided CNN with genetic algorithm. *Neural Comput Appl* 35(7):5165–5191
 49. Soydaner D (2022) Attention mechanism in neural networks: where it comes and where it goes. *Neural Comput Appl* 34(16):13371–13385
 50. Li W, Fan L, Wang Z, Ma C, Cui X (2021) Tackling mode collapse in multi-generator GANs with orthogonal vectors. *Pattern Recognit* 110:107646
 51. Zhang R, Isola P, Efros AA, Shechtman E, Wang O (2018) The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 586–595
 52. Johnson J, Alahi A, Fei-Fei L (2016) Perceptual losses for real-time style transfer and super-resolution. In: European Conference on Computer Vision, pp. 694–711. Springer
 53. Zhang Z, Schomaker L (2022) Optimized latent-code selection for explainable conditional text-to-image gans. In: 2022 International Joint Conference on Neural Networks (IJCNN), pp. 1–9. IEEE
 54. Weinberger KQ, Saul LK (2006) Unsupervised learning of image manifolds by semidefinite programming. *Int J Comput Vis* 70(1):77–90
 55. Cadoni M, Lagorio A, Khellat-Kihel S, Grosso E (2021) On the correlation between human fixations, handcrafted and CNN features. *Neural Comput Appl* 33:11905–11922
 56. Bengio Y, Mesnil G, Dauphin Y, Rifai S (2013) Better mixing via deep representations. In: International Conference on Machine Learning, pp. 552–560. PMLR
 57. Cui J (2014) 2d and 3d palmprint fusion and recognition using PCA plus TPTSR method. *Neural Comput Appl* 24:497–502
 58. Deng W, Liu Y, Hu J, Guo J (2012) The small sample size problem of ica: A comparative study and analysis. *Pattern Recognit* 45(12):4438–4450
 59. Wen L, Li X, Gao L (2020) A transfer convolutional neural network for fault diagnosis based on resnet-50. *Neural Comput Appl* 32:6111–6124
 60. Zagoruyko S, Komodakis N (2016) Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. arXiv preprint [arXiv:1612.03928](https://arxiv.org/abs/1612.03928)
 61. Smith LN (2017) Cyclical learning rates for training neural networks. In: 2017 IEEE Winter Conference on Applications of Computer Vision (WACV), pp. 464–472. IEEE
 62. Oosten J-P, Schomaker L (2014) Separability versus prototypicality in handwritten word-image retrieval. *Pattern Recognit* 47(3):1031–1038
 63. Leporowski B, Iosifidis A (2021) Visualising deep network time-series representations. *Neural Comput Appl* 33:16489–16498
 64. Ijjina EP, Chalavadi KM (2017) Human action recognition in RGB-d videos using motion sequence information and deep learning. *Pattern Recognit* 72:504–516
 65. Jiang P-T, Zhang C-B, Hou Q, Cheng M-M, Wei Y (2021) Layercam: exploring hierarchical class activation maps. *IEEE Trans Image Process* 38:5875–5888
 66. Sundararajan M, Taly A, Yan Q (2017) Axiomatic attribution for deep networks. In: International Conference on Machine Learning, pp. 3319–3328. PMLR
 67. Fong R, Patrick M, Vedaldi A (2019) Understanding deep networks via extremal perturbations and smooth masks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 2950–2958

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.