

# A survey and taxonomy of adversarial neural networks for text-to-image synthesis

Jorge Agnese<sup>1</sup> | Jonathan Herrera<sup>1</sup> | Haicheng Tao<sup>2</sup> | Xingquan Zhu<sup>1</sup> 

<sup>1</sup>Department of Computer & Electrical Engineering and Computer Science, Florida Atlantic University, Boca Raton, Florida

<sup>2</sup>Jiangsu Provincial Key Laboratory of E-Business, Nanjing University of Finance and Economics, Nanjing, China

## Correspondence

Xingquan Zhu, Department of Computer & Electrical Engineering and Computer Science, Florida Atlantic University, Boca Raton, FL 33431.  
Email: xzhu3@fau.edu

## Funding information

National Natural Science Foundation of China (NSFC), Grant/Award Number: 91646204; US National Science Foundation (NSF), Grant/Award Numbers: CNS-1828181, IIS-1763452

## Abstract

Text-to-image synthesis refers to computational methods which translate human written textual descriptions, in the form of keywords or sentences, into images with similar semantic meaning to the text. In earlier research, image synthesis relied mainly on word to image correlation analysis combined with supervised methods to find best alignment of the visual content matching to the text. Recent progress in deep learning (DL) has brought a new set of unsupervised DL methods, particularly deep generative models which are able to generate realistic visual images using suitably trained neural network models. The change of direction from the computer vision-based approaches to artificial intelligence (AI)-driven methods ignited the intense interest in industry, such as virtual reality, recreational & professional (eSports) gaming, and computer-aided design, and so on, to automatically generate compelling images from text-based natural language descriptions. In this paper, we review the most recent development in the text-to-image synthesis research domain. Our goal is to provide value by delivering a comparative review of the state-of-the-art models in terms of their architecture and design. The survey first introduces image synthesis and its challenges, and then reviews key concepts such as generative adversarial networks (GANs) and deep convolutional encoder-decoder neural networks (DCNNs). After that, we propose a taxonomy to summarize GAN-based text-to-image synthesis into four major categories: semantic enhancement GANs, resolution enhancement GANs, diversity enhancement GANs, and motion enhancement GANs. We elaborate on the main objective of each group, and further review typical GAN architectures in each group. The taxonomy and the review outline the techniques and the evolution of different approaches, and eventually provide a roadmap to summarize the list of contemporaneous solutions that utilize GANs and DCNNs to generate enthralling results in categories such as human faces, birds, flowers, room interiors, object reconstruction from edge maps (games), and so on. The survey

**Abbreviations:** AI, artificial intelligence; DCNN, deep convolutional encoder-decoder neural network; DL, deep learning; GAN, generative adversarial network.

Jorge Agnese, Jonathan Herrera, and Haicheng Tao contributed equally to this study.

concludes with a comparison of the proposed solutions, challenges that remain unresolved, and future developments in the text-to-image synthesis domain.

This article is categorized under:

Algorithmic Development > Multimedia

Technologies > Machine Learning

#### KEYWORDS

deep learning, generative adversarial network (GAN), machine learning, text-to-image synthesis

## 1 | INTRODUCTION

(GANs), and the variations that are now being proposed is the most interesting idea in the last 10 years in ML, in my opinion. (2016)  
– Yann LeCun

A picture is worth a thousand words! While written text provides efficient, effective, and concise ways for communication, visual content, such as images, is a more comprehensive, accurate, and intelligible method of information sharing and understanding. Generation of images from text descriptions, that is, text-to-image synthesis, is a complex computer vision and machine learning problem that has seen great progress over recent years. Automatic image generation from natural language may allow users to describe visual elements through visually rich text descriptions. The ability to do so effectively is highly desirable as it could be used in artificial intelligence (AI) applications such as computer-aided design, image editing (Chen, Shen, Gao, Liu, & Liu, 2018; Yan, Zhang, Wang, Paris, & Yu, 2016), game engines for the development of the next generation of video games (Isola, Zhu, Zhou, & Efros, 2018), and pictorial art generation (Elgammal, Liu, Elhoseiny, & Mazzone, 2017).

### 1.1 | Traditional learning-based text-to-image synthesis


In the early stages of research, text-to-image synthesis was mainly carried out through a search and supervised learning combined process (Zhu, Goldberg, Eldawy, Dyer, & Strock, 2007), as shown in Figure 1. In order to connect text descriptions to images, one could use correlation between keywords (or keyphrase) and images that identifies informative and “picturable” text units; then, these units would search for the most likely image parts conditioned on the text, eventually optimizing the picture layout conditioned on both the text and the image parts. Such methods often integrated multiple AI key components, including natural language processing, computer vision, computer graphics, and machine learning.

The major limitation of the traditional learning-based text-to-image synthesis approaches is that they lack the ability to generate new image content; they can only change the characteristics of the given/training images. Alternatively, research in generative models has advanced significantly and delivers solutions to learn from training images and produce new visual content. For example, Attribute2Image (Yan, Yang, Sohn, & Lee, 2016) models each image as a composite of foreground and background. In addition, a layered generative model with disentangled latent variables is learned, using a variational auto-encoder, to generate visual content. Because the learning is customized/conditioned by given attributes, the generative models of Attribute2Image can generate images with respect to different attributes, such as gender, hair color, age, and so on, as shown in Figure 2.

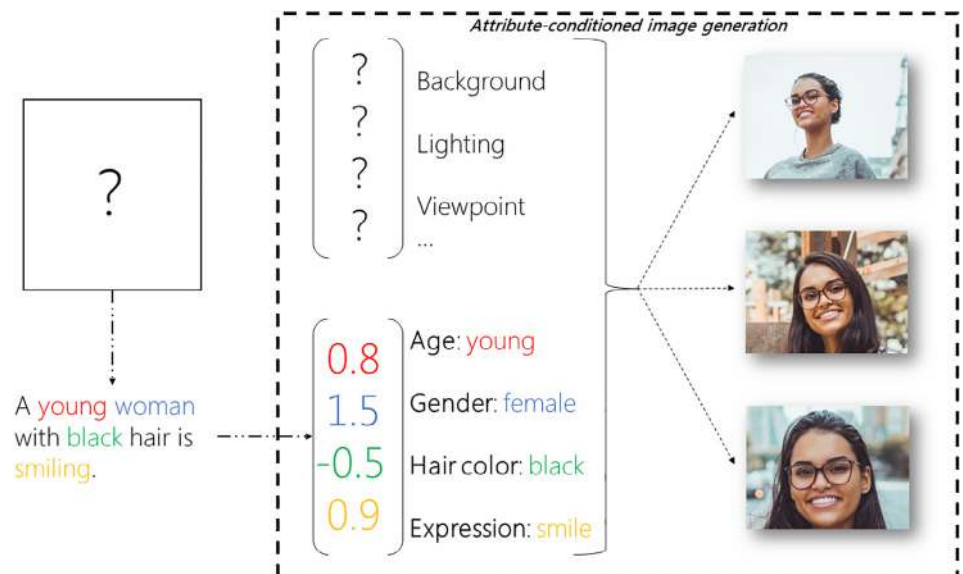
### 1.2 | GAN-based text-to-image synthesis

Although generative model-based text-to-image synthesis provides much more realistic image synthesis results, the image generation is still conditioned by the limited attributes. In recent years, several papers have been published on the subject of text-to-image synthesis. Most of the contributions from these papers rely on multimodal learning

**FIGURE 1** Early research on text-to-image synthesis (Zhu et al., 2007). The system uses correlation between keywords (or keyphrase) and images and identifies informative and “picturable” text units, then searches for the most likely image parts conditioned on the text, and eventually optimizes the picture layout conditioned on both the text and image parts



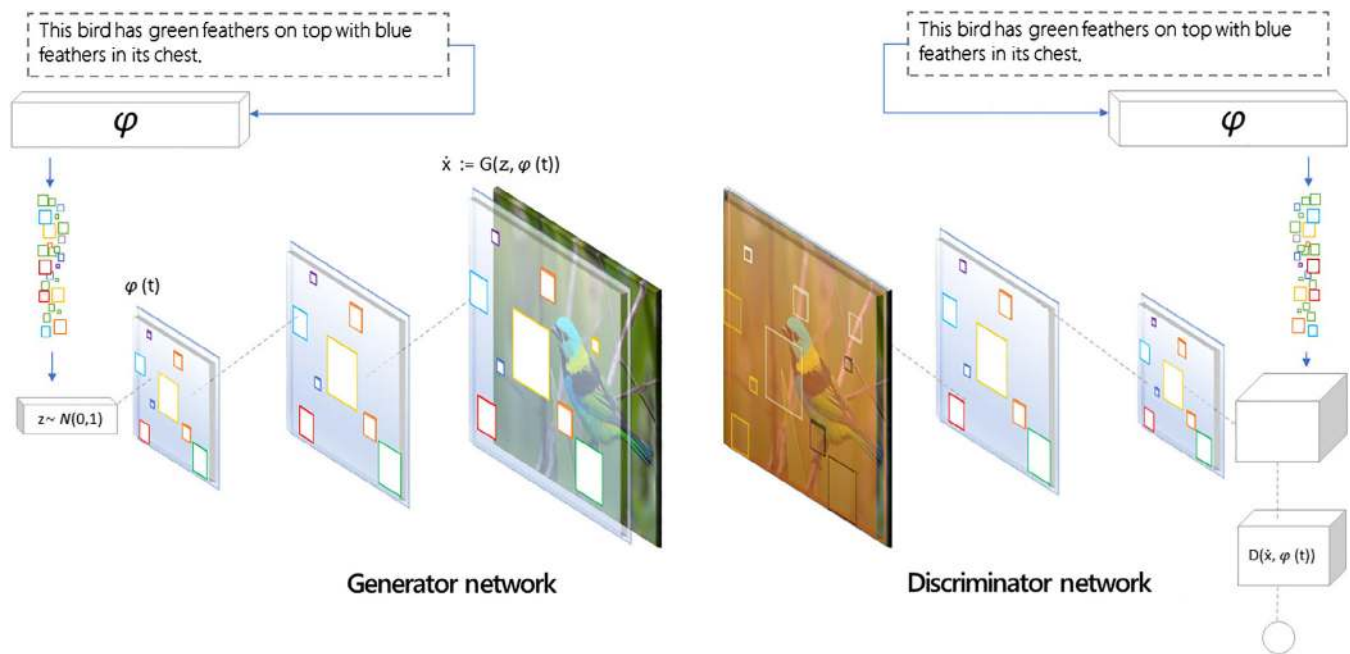
**FIGURE 2** Supervised learning-based text-to-image synthesis (Yan, Yang, et al., 2016). The supervised learning process aims to learn layered generative models to generate visual content. Because the learning is customized/conditioned by the given attributes, the generative models of Attribute2Image can generate images with respect to different attributes, such as hair color, age, and so on



approaches that include generative adversarial networks (GANs) and deep convolutional decoder networks as their main drivers to generate entrancing images from text (Goodfellow et al., 2014; Odena, Olah, & Shlens, 2017; Reed et al., 2016; Wu, Xu, & Hall, 2017; Xu et al., 2018).

First introduced by Ian Goodfellow et al. (2014), GANs consist of two neural networks paired with a discriminator and a generator. These two models compete with one another, with the generator attempting to produce synthetic/fake samples that will fool the discriminator and the discriminator attempting to differentiate between real (genuine) and synthetic samples. Because GANs' adversarial training aims to cause generators to produce images similar to the real (training) images, GANs can naturally be used to generate synthetic images (image synthesis), and this process can even be customized further by using text descriptions to specify the types of images to generate, as shown in Figure 3.

Much like text-to-speech and speech-to-text conversion, there exists a wide variety of problems that text-to-image synthesis could solve in the computer vision field specifically (Haynes, Norton, McParland, & Cooper, 2018; Reed, Akata, Yan, et al., 2016). Nowadays, researchers are attempting to solve a plethora of computer vision problems with the aid of deep convolutional networks, GANs, and a combination of multiple methods, often called multimodal



**FIGURE 3** Generative adversarial network (GAN)-based text-to-image synthesis (Huang, Yu, & Wang, 2018). GAN-based text-to-image synthesis combines discriminative and generative learning to train neural networks resulting in the generated images semantically resemble to the training samples or tailored to a subset of training images (i.e., conditioned outputs).  $\varphi()$  is a feature embedding function, which converts text as feature vector.  $z$  is a latent vector following normal distributions with zero mean.  $\hat{x} = G(z, \varphi(t))$  denotes a synthetic image generated from the generator, using latent vector  $z$  and the text features  $\varphi(t)$  as the input.  $D(\hat{x}, \varphi(t))$

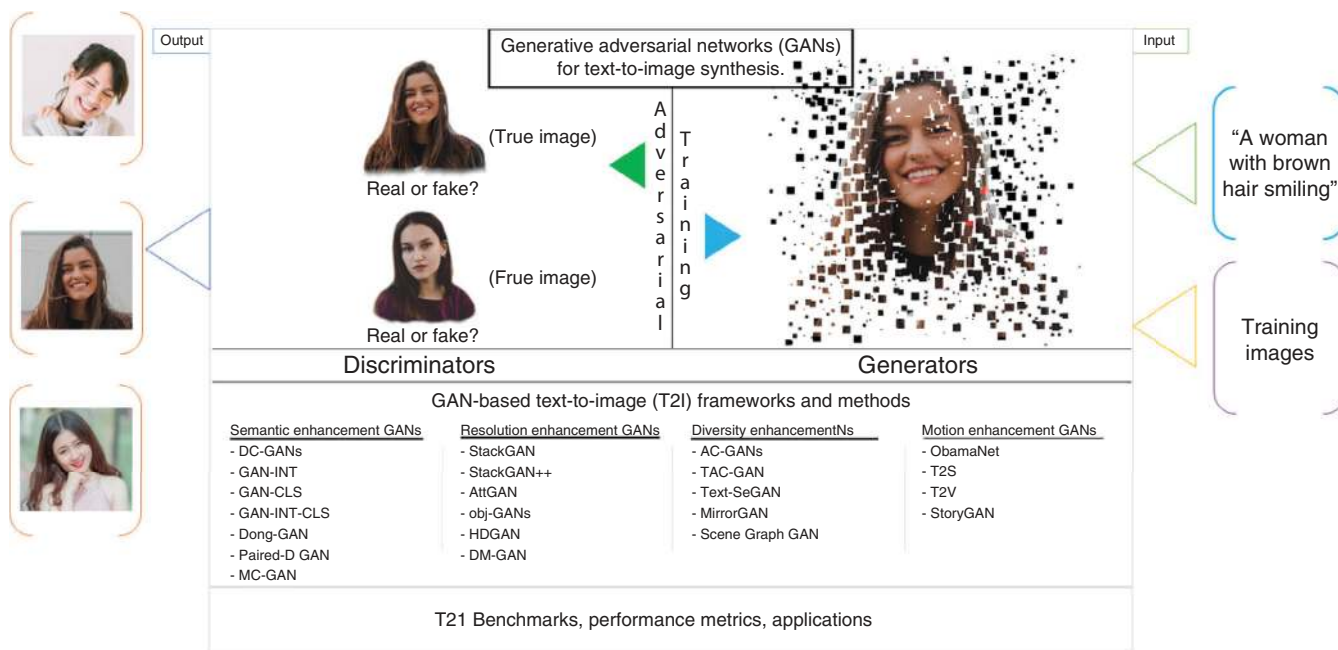
learning methods (Reed, Akata, Yan, et al., 2016). For simplicity, *multiple learning methods* will be referred to as *multimodal learning* hereafter (Ngiam et al., 2011). Researchers often describe multimodal learning as a method that incorporates characteristics coming from several methods, algorithms, and ideas. This can include ideas from two or more learning approaches in order to create a robust implementation to solve an uncommon problem or improve a solution (Baltrusaitis, Ahuja, & Morency, 2017; Dash, Gamboa, Ahmed, Afzal, & Liwicki, 2017; Li, Wang, & Qi, 2018; Reed, Akata, Yan, et al., 2016; Yang, Kannan, Batra, & Parikh, 2017).

In this survey, we focus primarily on reviewing recent works that aim to solve the challenge of text-to-image synthesis using GANs. In order to provide a clear roadmap, we propose a taxonomy to summarize reviewed GANs into four major categories. Our review will elaborate the motivations of methods in each category, analyze typical models, their network architectures, and possible drawbacks for further improvement. The visual abstract of the survey and the list of reviewed GAN frameworks are shown in Figure 4.

The remainder of the survey is organized as follows. Section 2 presents a brief summary of existing works on subjects similar to that of this paper and highlights the key distinctions making ours unique. Section 3 gives a short introduction to GANs and some preliminary concepts related to image generation, as they are the engines that make text-to-image synthesis possible and are essential building blocks to achieve photo-realistic images from text descriptions. Section 4 proposes a taxonomy to summarize GAN-based text-to-image synthesis, discusses models and architectures of novel works focused solely on text-to-image synthesis. This section will also draw key contributions from these works in relation to their applications. Section 5 reviews GAN-based text-to-image synthesis benchmarks, performance metrics, and comparisons, including a simple review of GANs for other applications. In Section 6, we conclude with a brief summary and outline ideas for future interesting developments in the field of text-to-image synthesis.

## 2 | RELATED WORK

With the growth and success of GANs, deep convolutional decoder networks, and multimodal learning methods, these techniques were some of the first procedures which aimed to solve the challenge of image synthesis. Many engineers and scientists in computer vision and AI have contributed through extensive studies and experiments, with numerous



**FIGURE 4** A visual summary of generative adversarial network (GAN)-based text-to-image (T2I) synthesis process, and the summary of GAN-based frameworks/methods reviewed in the survey

proposals and publications detailing their contributions. Because GANs, introduced by Goodfellow et al. (2014), are emerging research topics, their practical applications to image synthesis are still in their infancy. Recently, many new GAN architectures and designs have been proposed to use GANs for different applications, for example, using GANs to generate sentimental texts (Wang & Wan, 2018), or using GANs to transform natural images into cartoons (Chen, Lai, & Liu, 2018).

Although GANs are becoming increasingly popular, very few survey papers currently exist to summarize and outline contemporaneous technical innovations and contributions of different GAN architectures (Creswell et al., 2018; Hong, Hwang, Yoo, & Yoon, 2019). Survey papers specifically attuned to analyzing different contributions to text-to-image synthesis using GANs are even more scarce. We have thus found two surveys (Huang et al., 2018; Wu et al., 2017) on image synthesis using GANs, which are the two most closely related publications to our survey objective. In the following paragraphs, we briefly summarize each of these surveys and point out how our objectives differ from theirs.

In Huang et al. (2018), the authors provide an overview of image synthesis using GANs. In this survey, the authors discuss the motivations for research on image synthesis and introduce some background information on the history of GANs, including a section dedicated to core concepts of GANs, namely generators, discriminators, and the min-max game analogy, and some enhancements to the original GAN model, such as conditional GANs, addition of variational autoencoders, and so on. In this survey, we will carry out a similar review of the background knowledge because the understanding of these preliminary concepts is paramount for the rest of the paper. Three types of approaches for image generation are reviewed, including direct methods (single generator and discriminator), hierarchical methods (two or more generator-discriminator pairs, each with a different goal), and iterative methods (each generator-discriminator pair generates a gradually higher resolution image). Following the introduction, Huang et al. (2018) discusses methods for text-to-image and image-to-image synthesis, respectively, and also describes several evaluation metrics for synthetic images, including inception scores (ISs) and Frechet inception distance (FID), and explains the significance of the discriminators acting as learned loss functions as opposed to fixed loss functions.

Different from the above survey, which has a relatively broad scope in GANs, our objective is heavily focused on text-to-image synthesis. Although this topic, text-to-image synthesis, has indeed been covered in Huang et al. (2018), they did so in a much less detailed fashion, mostly listing the many different works in a time-sequential order. In comparison, we will review several representative methods in the field and outline their models and contributions in detail.

Similarly to Huang et al. (2018), the second survey paper (Wu et al., 2017) begins with a standard introduction addressing the motivation of image synthesis and the challenges it presents followed by a section dedicated to core

concepts of GANs and enhancements to the original GAN model. In addition, the paper covers the review of two types of applications: (1) unconstrained applications of image synthesis such as super-resolution, image inpainting, and so on, and (2) constrained image synthesis applications, namely image-to-image, text-to-image, and sketch-to image, and also discusses image and video editing using GANs. Again, the scope of this paper is intrinsically comprehensive, while we focus specifically on text-to-image and go into more detail regarding the contributions of novel state-of-the-art models.

Other surveys have been published on related matters, mainly related to the advancements and applications of GANs (Zhang, Tu, & Cui, 2017; Zhang, Zhai, Luo, Zhan, & Chen, 2018), but we have not found any prior works which focus specifically on text-to-image synthesis using GANs. To our knowledge, this is the first paper to do so.

### 3 | PRELIMINARIES AND FRAMEWORKS

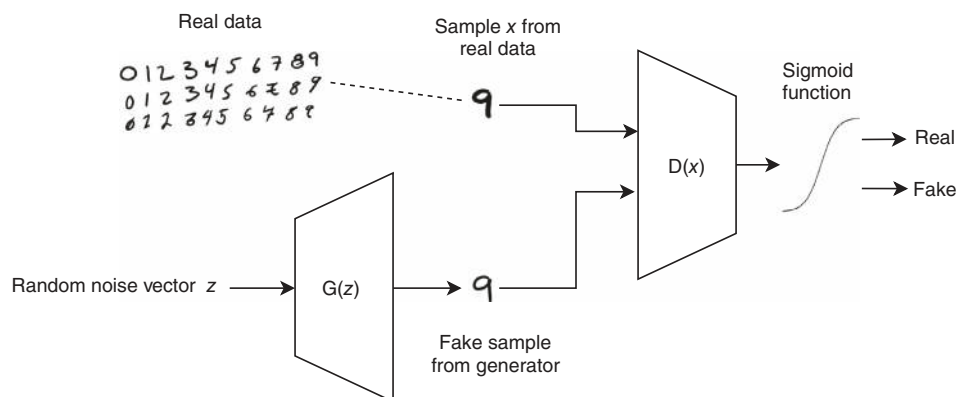
In this section, we first introduce preliminary knowledge of GANs and one of its commonly used variants, conditional GAN (i.e., conditional generative adversarial networks [cGAN]), which is the building block for many GAN-based text-to-image synthesis models. After that, we briefly separate GAN-based text-to-image synthesis into two types, Simple GAN frameworks versus Advanced GAN frameworks, and discuss why advanced GAN architecture for image synthesis.

Note that the simple versus advanced GAN framework separation is rather too brief, our taxonomy in the next section will propose a taxonomy to summarize advanced GAN frameworks into four categories, based on their objective and designs.

#### 3.1 | Generative adversarial neural network

Before moving on to a discussion and analysis of works applying GANs for text-to-image synthesis, there are some preliminary concepts, enhancements of GANs, datasets, and evaluation metrics that are present in some of the works described in the next section and are thus worth introducing.

As stated previously, GANs were introduced by Ian Goodfellow et al. (2014), and consist of two deep neural networks, a generator and a discriminator, which are trained independently with conflicting goals: The generator aims to generate samples closely related to the original data distribution and fool the discriminator, while the discriminator aims to distinguish between samples from the generator model and samples from the true data distribution by calculating the probability of the sample coming from either source. A conceptual view of the GAN architecture is shown in Figure 5.



**FIGURE 5** A conceptual view of the generative adversarial network (GAN) architecture. The Generator  $G(z)$  is trained to generate synthetic/fake resemble to real samples, from a random noise distribution. The fake samples are fed to the Discriminator  $D(x)$  along with real samples. The Discriminator is trained to differentiate fake samples from real samples. The iterative training of the generator and the discriminator helps GAN deliver good generator generating samples very close to the underlying training samples

The training of GANs is an iterative process that, with each iteration, updates the generator and the discriminator with the goal of each defeating the other, leading each model to become increasingly adept at its specific task until a threshold is reached. This is analogous to a min-max game between the two models, according to the following equation:

$$\min_{\theta_g} \max_{\theta_d} V(D_{\theta_d}, G_{\theta_g}) = \mathbb{E}_{x \sim P_{data}(x)} [\log(D_{\theta_d}(x))] + \mathbb{E}_{z \sim P_z(z)} [\log(1 - D_{\theta_d}(G_{\theta_g}(z)))] \quad (1)$$

In Equation (1),  $x$  denotes a multidimensional sample, for example, an image, and  $z$  denotes a multidimensional latent space vector, for example, a multidimensional data point following a predefined distribution function such as that of normal distributions.  $D_{\theta_d}()$  denotes a discriminator function, controlled by parameters  $\theta_d$ , which aims to classify a sample into a binary space.  $G_{\theta_g}()$  denotes a generator function, controlled by parameters  $\theta_g$ , which aims to generate a sample from some latent space vector. For example,  $G_{\theta_g}(z)$  means using a latent vector  $z$  to generate a synthetic/fake image, and  $D_{\theta_d}(x)$  means to classify an image  $x$  as binary output (i.e., true/false or 1/0). In the GAN setting, the discriminator  $D_{\theta_d}()$  is learned to distinguish a genuine/true image (labeled as 1) from fake images (labeled as 0). Therefore, given a true image  $x$ , the ideal output from the discriminator  $D_{\theta_d}(x)$  would be 1. Given a fake image generated from the generator  $G_{\theta_g}(z)$ , the ideal prediction from the discriminator  $D_{\theta_d}(G_{\theta_g}(z))$  would be 0, indicating the sample is a fake image.

Following the above definition, the minmax objective function in Equation (1) aims to learn parameters for the discriminator ( $\theta_d$ ) and generator ( $\theta_g$ ) to reach an optimization goal: The discriminator intends to differentiate true versus fake images with maximum capability  $\max_{\theta_d}$  whereas the generator intends to minimize the difference between a fake image versus a true image  $\min_{\theta_g}$ . In other words, the discriminator sets the characteristics and the generator produces elements, often images, iteratively until it meets the attributes set forth by the discriminator. GANs are often used with images and other visual elements and are notoriously efficient in generating compelling and convincing photorealistic images. Most recently, GANs were used to generate an original painting in an unsupervised fashion (Radford, Metz, & Chintala, 2016). The following sections go into further detail regarding how the generator and discriminator are trained in GANs.

### 3.1.1 | Generator

In image synthesis, the generator network can be thought of as a mapping from one representation space (latent space) to another (actual data) (Creswell et al., 2018). When it comes to image synthesis, all of the images in the data space fall into some distribution in a very complex and high-dimensional feature space. Sampling from such a complex space is very difficult, so GANs instead train a generator to create synthetic images from a much more simple feature space (usually random noise) called the latent space. The generator network performs up-sampling of the latent space and is usually a deep neural network consisting of several convolutional and/or fully connected layers (Creswell et al., 2018). The generator is trained using gradient descent to update the weights of the generator network with the aim of producing data (in our case, images) that the discriminator classifies as real.

### 3.1.2 | Discriminator

The discriminator network can be thought of as a mapping from image data to the probability of the image coming from the real data space, and is also generally a deep neural network consisting of several convolution and/or fully connected layers. However, the discriminator performs down-sampling as opposed to up-sampling. Like the generator, it is trained using gradient descent but its goal is to update the weights so that it is more likely to correctly classify images as real or fake.

In GANs, the ideal outcome is for both the generator's and discriminator's cost functions to converge so that the generator produces photorealistic images that are indistinguishable from real data, and the discriminator at the same time becomes an expert at differentiating between real and synthetic data. This, however, is not possible since a reduction in cost of one model generally leads to an increase in cost of the other. This phenomenon makes training GANs very difficult, and training them simultaneously (both models performing gradient descent in parallel) often leads to a

stable orbit where neither model are able to converge. To combat this, the generator and discriminator are often trained independently. In this case, the GAN remains the same, but there are different training stages. In one stage, the weights of the generator are kept constant and gradient descent updates the weights of the discriminator, and in the other stage the weights of the discriminator are kept constant while gradient descent updates the weights of the generator. This is repeated for some number of epochs until a desired low cost for each model is reached (Salimans et al., 2016).

### 3.2 | cGAN: Conditional GAN

Conditional generative adversarial networks (cGAN) are an enhancement of GANs proposed by Mirza and Osindero (2014) shortly after the introduction of GANs by Goodfellow et al. (2014). The objective function of the cGAN is defined in Equation (2) which is very similar to the GAN objective function in Equation (1) except that the inputs to both discriminator and generator are conditioned by a class label  $y$ .

$$\min_{\theta_g} \max_{\theta_d} V(D_{\theta_d}, G_{\theta_g}) = \mathbb{E}_{x \sim P_{data}(x)} [\log(D_{\theta_d}(x|y))] + \mathbb{E}_{z \sim P_z(z)} [\log(1 - D_{\theta_d}(G_{\theta_g}(z|y)))] \quad (2)$$

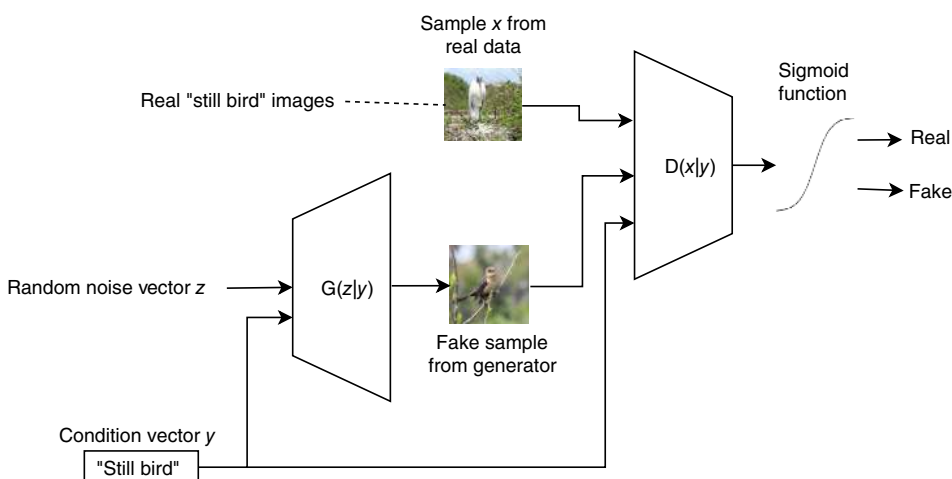
The main technical innovation of cGAN is that it introduces an additional input or inputs to the original GAN model, allowing the model to be trained on information such as class labels or other conditioning variables as well as the samples themselves, concurrently. Whereas the original GAN was trained only with samples from the data distribution, resulting in the generated sample reflecting the general data distribution, cGAN enables directing the model to generate more tailored outputs.

In Figure 6, the condition vector is the class label (text string) “Red bird,” which is fed to both the generator and discriminator. It is important, however, that the condition vector is related to the real data. If the model in Figure 6 was trained with the same set of real data (red birds) but the condition text was “Yellow fish,” the generator would learn to create images of red birds when conditioned with the text “Yellow fish.”

Note that the condition vector in cGAN can come in many forms, such as texts, not just limited to the class label. Such a unique design provides a direct solution to generate images conditioned by predefined specifications. As a result, cGAN has been used in text-to-image synthesis since the very first day of its invention although modern approaches can deliver much better text-to-image synthesis results.

### 3.3 | Simple GAN frameworks for text-to-image synthesis

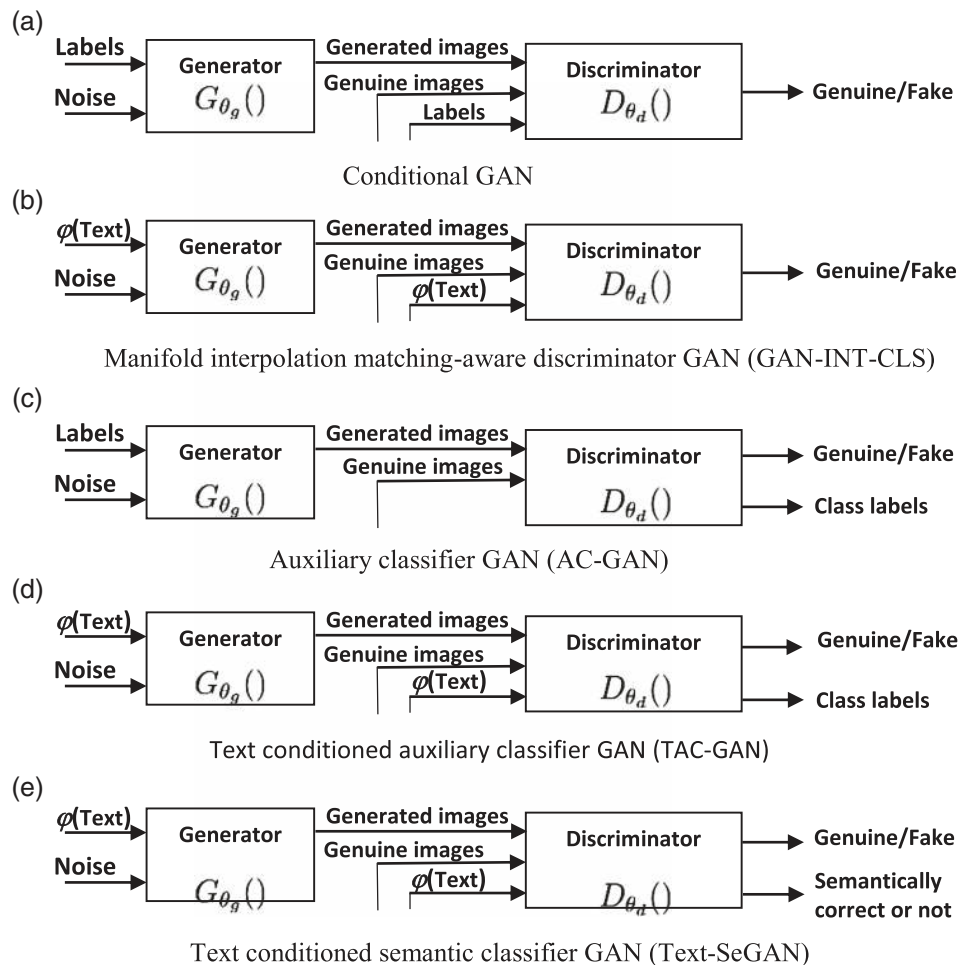
In order to generate images from text, one simple solution is to employ the conditional GAN (cGAN) designs and add conditions to the training samples, such that the GAN is trained with respect to the underlying conditions. Several pioneer works have followed similar designs for text-to-image synthesis.



**FIGURE 6** A conceptual view of the conditional GAN (cGAN) architecture. The Generator  $G(z|y)$  generates samples from a random noise distribution and some condition vector (in this case text). The fake samples are fed to the Discriminator  $D(x|y)$  along with real samples and the same condition vector, and the Discriminator calculates the probability that the fake sample came from the real data distribution



An essential disadvantage of using cGAN for text-to-image synthesis is that it cannot handle complicated textual descriptions for image generation, because cGAN uses labels as conditions to restrict the GAN inputs. If the text inputs have multiple keywords (or long-text descriptions) they cannot be used simultaneously to restrict the input. Instead of using text as conditions, another two approaches (Dash, Gamboa, Ahmed, Afzal, & Liwicki, 2017; Reed, Akata, Yan, et al., 2016) use text as input features, and concatenate such features with other features to train discriminator and generator, as shown in Figure 7b and c. To ensure text being used as GAN input, a feature embedding or feature representation learning (Bengio, Courville, & Vincent, 2013; Zhang, Yin, Zhu, & Zhang, 2018) function  $\varphi()$  is often introduced to convert input text as numeric features, which are further concatenated with other features to train GANs.



**FIGURE 7** A simple architecture comparisons between five generative adversarial network (GAN) networks for text-to-image synthesis. This figure also explains how texts are fed as input to train GAN to generate images. (a) Conditional GAN (cGAN) (Mirza & Osindero, 2014) use labels to condition the input to the generator and the discriminator. The final output is discriminator similar to generic GAN; (b) Manifold interpolation matching-aware discriminator GAN (GAN-INT-CLS) (Reed, Akata, Yan, et al., 2016) feeds text input to both generator and discriminator (texts are preprocessed as embedding features, using function  $\varphi()$ , and concatenated with other input, before feeding to both generator and discriminator). The final output is discriminator similar to generic GAN; (c) Auxiliary classifier GAN (AC-GAN) (Odena, Olah, & Shlens, 2017) uses an auxiliary classifier layer to predict the class of the image to ensure that the output consists of images from different classes, resulting in diversified synthesis images; (d) text conditioned auxiliary classifier GAN (TAC-GAN) (Dash, Gamboa, Ahmed, Afzal, & Liwicki, 2017) share similar design as GAN-INT-CLS, whereas the output include both a discriminator and a classifier (which can be used for classification); and (e) text conditioned semantic classifier GAN (Text-SeGAN) (Cha, Gwon, & Kung, 2019) uses a regression layer to estimate the semantic relevance between the image, so the generated images are not limited to certain classes and are semantically matching to the text input

### 3.4 | Advanced GAN frameworks for text-to-image synthesis

Motivated by the GAN and conditional GAN (cGAN) design, many GAN-based frameworks have been proposed to generate images, with different designs and architectures, such as using multiple discriminators, using progressively trained discriminators, or using hierarchical discriminators. Figure 8 outlines several advanced GAN frameworks in the literature. In addition to these frameworks, many new designs are being proposed to advance the field with rather sophisticated designs. For example, a recent work (Gao et al., 2019) proposes to use a pyramid generator and three independent discriminators, each focusing on a different aspect of the images, to lead the generator toward creating images that are photorealistic on multiple levels. Another recent publication (Cha, Gwon, & Kung, 2019) proposes to use discriminator to measure semantic relevance between image and text instead of class prediction (like most discriminator in GANs does), resulting a new GAN structure outperforming text conditioned auxiliary classifier (TAC-GAN) (Dash, Gamboa, Ahmed, Afzal, & Liwicki, 2017) and generating diverse, realistic, and relevant to the input text regardless of class.

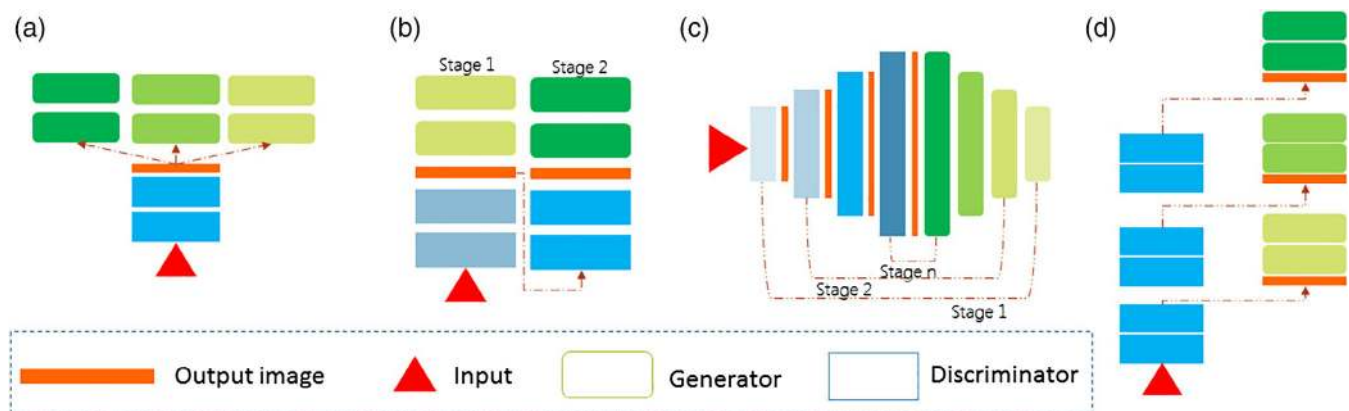
In the following section, we will first propose a taxonomy that summarizes advanced GAN frameworks for text-to-image synthesis, and review most recent proposed solutions to the challenge of generating photorealistic images conditioned on natural language text descriptions using GANs. The solutions we discuss are selected based on relevance and quality of contributions. Many publications exist on the subject of image generation using GANs, but in this paper, we focus specifically on models for text-to-image synthesis, with the review emphasizing on the “model” and “contributions” for text-to-image synthesis. At the end of this section, we also briefly review methods using GANs for other image-synthesis applications.

## 4 | TEXT-TO-IMAGE SYNTHESIS TAXONOMY AND CATEGORIZATION

In this section, we propose a taxonomy to summarize advanced GAN-based text-to-image synthesis frameworks, as shown in Figure 9. The taxonomy organizes GAN frameworks into four categories, including semantic enhancement GANs, resolution enhancement GANs, diversity enhancement GANs, and motion enhancement GAGs. Following the proposed taxonomy, each subsection will introduce several typical frameworks and address their techniques of using GANs to solve certain aspects of the text-to-image synthesis challenges.

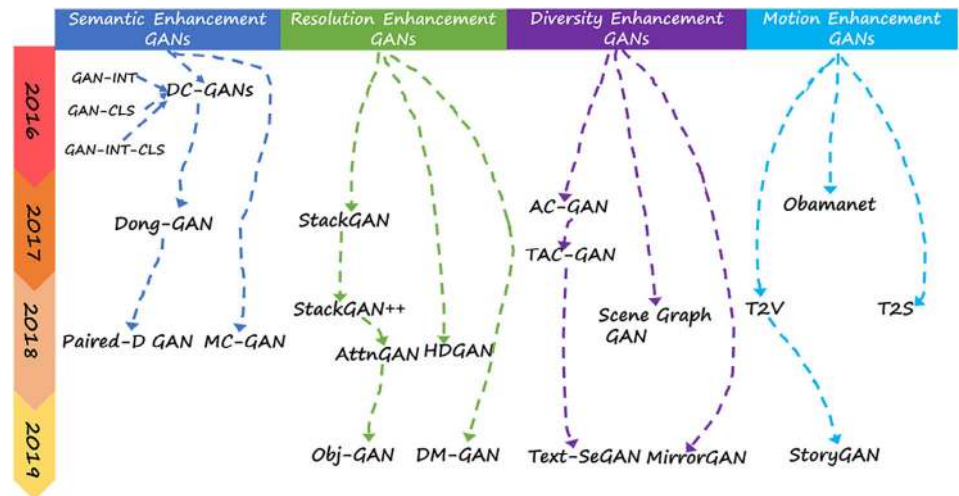
### 4.1 | GAN-based text-to-image synthesis taxonomy

Although the ultimate goal of Text-to-Image synthesis is to generate images closely related to the textual descriptions, the relevance of the images to the texts are often validated from different perspectives, due to the inherent diversity of



**FIGURE 8** A high level comparison of several advanced generative adversarial networks (GANs) framework for text-to-image synthesis. All frameworks take text (red triangle) as input and generate output images. From left to right, (a) uses multiple discriminators and one generator (Durugkar, Gemp, & Mahadevan, 2017; Nguyen, Le, Vu, & Phung, 2017), (b) uses multiple stage GANs where the output from one GAN is fed to the next GAN as input (Denton, Chintala, szlam, & Fergus, 2015; Zhang et al., 2017), (c) progressively trains symmetric discriminators and generators (Huang, Li, Poursaeed, Hopcroft, & Belongie, 2017), and (d) uses a single-stream generator with a hierarchically nested discriminator trained from end-to-end (Zhang, Xie, & Yang, 2018)

**FIGURE 9** A Taxonomy and categorization of advanced generative adversarial network (GAN) frameworks for text-to-image synthesis. We categorize advanced GAN frameworks into four major categories: semantic enhancement GANs, resolution enhancement GANs, diversity enhancement GANs, and motion enhancement GAGs. The relationship between relevant frameworks and their publication date are also outlined as a reference



human perceptions. For example, when generating images matching to the description “rose flowers,” some users may know the exact type of flowers they like and intend to generate rose flowers with similar colors. Other users may seek to generate high quality rose flowers with a nice background (e.g., garden). The third group of users may be more interested in generating flowers similar to rose but with different colors and visual appearance, for example, roses, begonia, and peony. The fourth group of users may want to not only generate flower images, but also use them to form a meaningful action, for example, a video clip showing flower growth, performing a magic show using those flowers, or telling a love story using the flowers.

From the text-to-image synthesis point of view, the first group of users intends to precisely control the semantic of the generated images, and their goal is to match the texts and images at the semantic level. The second group of users is more focused on the resolutions and the quality of the images, in addition to the requirement that the images and texts are semantically related. For the third group of users, their goal is to diversify the output images, such that their images carry diversified visual appearances and are also semantically related. The fourth user group adds a new dimension in image synthesis, and aims to generate sequences of images which are coherent in temporal order, that is, capture the motion information.

Based on the above descriptions, we categorize GAN-based text-to-image synthesis into a taxonomy with four major categories, as shown in Figure 9.

- *Semantic enhancement GANs*: Semantic enhancement GANs represent pioneer works of GAN frameworks for text-to-image synthesis. The main focus of the GAN frameworks is to ensure that the generated images are semantically related to the input texts. This objective is mainly achieved by using a neural network to encode texts as dense features, which are further fed to a second network to generate images matching to the texts.
- *Resolution enhancement GANs*: Resolution enhancement GANs mainly focus on generating high-quality images which are semantically matched to the texts. This is mainly achieved through a multistage GAN framework, where the outputs from earlier stage GANs are fed to the second (or later) stage GAN to generate better quality images.
- *Diversity enhancement GANs*: Diversity enhancement GANs intend to diversify the output images, such that the generated images are not only semantically related but also have different types and visual appearance. This objective is mainly achieved through an additional component to estimate semantic relevance between generated images and texts, in order to maximize the output diversity.
- *Motion enhancement GANs*: Motion enhancement GANs intend to add a temporal dimension to the output images, such that they can form meaningful actions with respect to the text descriptions. This goal is mainly achieved through a two-step process which first generates images matching to the “actions” of the texts, followed by a mapping or alignment procedure to ensure that images are coherent in the temporal order.

In the following, we will introduce how these GAN frameworks evolve for text-to-image synthesis, and will also review some typical methods of each category.

## 4.2 | Semantic enhancement GANs

Semantic relevance is one of the most important criteria of the text-to-image synthesis. For most GANs discussed in this survey, they are required to generate images semantically related to the text descriptions. However, the semantic relevance is a rather subjective measure, and images are inherently rich in terms of its semantics and interpretations. Therefore, many GANs are further proposed to enhance the text-to-image synthesis from different perspectives. In this subsection, we will review several classical approaches which are commonly served as text-to-image synthesis baseline.

### 4.2.1 | Deep convolution generative adversarial network (DC-GAN)

Deep convolution generative adversarial network (DC-GAN) (Reed, Akata, Yan, et al., 2016) represents the pioneer work for text-to-image synthesis using GANs. Its main goal is to train a DC-GAN on text features. During this process, these text features are encoded by another neural network. This neural network is a hybrid convolutional recurrent network at the character level. Concurrently, both neural networks have also feed-forward inference in the way they condition text features. Generating realistic images automatically from natural language text is the motivation of several of the works proposed in this computer vision field. However, actual AI systems are far from achieving this task (Li, Su, & Liu, 2018; Liu, Meng, Xiang, & Pan, 2018; Mirza & Osindero, 2014; Reed, Akata, Yan, et al., 2016; Wang & Gupta, 2016; Yang et al., 2019; Zhang, Tu, & Cui, 2017). Lately, recurrent neural networks (RNNs) led the way to develop frameworks that learn discriminatively on text features. At the same time, GANs began recently to show some promise on generating compelling images of a whole host of elements including but not limited to faces, birds, flowers, and non-common images such as room interiors (Reed, Akata, Yan, et al., 2016). DC-GAN is a multimodal learning model that attempts to bridge together both of the above mentioned unsupervised machine learning algorithms, the RNN and GANs, with the sole purpose of speeding the generation of text-to-image synthesis.

Deep learning (DL) shed some light to some of the most sophisticated advances in natural language representation, image synthesis (Huang et al., 2017; Reed, Akata, Yan, et al., 2016; Wang et al., 2017; Wu et al., 2017), and classification of generic data (Han, Li, & Zhu, 2019). However, a bulk of the latest breakthroughs in DL and computer vision was related to supervised learning (Reed, Akata, Yan, et al., 2016). Even though natural language and image synthesis were part of several contributions on the supervised side of DL, unsupervised learning saw recently a tremendous rise in input from the research community specially on two subproblems: text-based natural language and image synthesis (Cha, Gwon, & Kung, 2017; Dong, Zhang, McIlwraith, & Guo, 2017; Reed, Akata, Yan, et al., 2016; Yang et al., 2017; Zhang et al., 2018). These subproblems are typically subdivided as focused research areas. DC-GAN's contributions are mainly driven by these two research areas. In order to generate plausible images from natural language, DC-GAN contributions revolve around developing a straightforward yet effective GAN architecture and training strategy that allows natural text to image synthesis. These contributions are primarily tested on the Caltech-UCSD Birds and Oxford-102 Flowers datasets. Each image in these datasets carries five text descriptions. These text descriptions were created by the research team when setting up the evaluation environment. The DC-GANs model is subsequently trained on several subcategories. Subcategories in this research represent the training and testing subdatasets. The performance shown by these experiments displays a promising yet effective way to generate images from textual natural language descriptions (Reed, Akata, Yan, et al., 2016).

### 4.2.2 | DC-GAN extensions

Following the pioneer DC-GAN framework (Reed, Akata, Yan, et al., 2016), many researches propose revised network structures (e.g., different discriminators) in order to improve images with better semantic relevance to the texts. Based on the DC-GAN architecture, GAN-CLS with image-text matching discriminator, GAN-INT learned with text manifold interpolation and GAN-INT-CLS which combines both are proposed to find semantic match between text and image. Similar to the DC-GAN architecture, an adaptive loss function (i.e., perceptual loss, Johnson, Alahi, & Fei-Fei, 2016) is proposed for semantic image synthesis which can synthesize a realistic image that not only matches the target text description but also keep the irrelevant features (e.g., background) from source images (Dong, Yu, Wu, & Guo, 2017). Regarding to the perceptual losses, three loss functions (i.e., pixel reconstruction loss, activation reconstruction loss, and texture reconstruction loss) are proposed in Cha et al. (2017) in which they construct the network architectures

based on the DC-GAN, that is, GAN-INT-CLS-Pixel, GAN-INT-CLS-VGG, and GAN-INT-CLS-Gram with respect to three losses. In Dong, Yu, et al. (2017), a residual transformation unit is added in the network to retain similar structure of the source image.

Following the Dong, Yu, et al. (2017) and considering the features in early layers address background while foreground is obtained in latter layers in CNN, a pair of discriminators with different architectures (i.e., Paired-D GAN) is proposed to synthesize background and foreground from a source image separately (Vo & Sugimoto, 2018). Meanwhile, the skip connection in the generator is employed to more precisely retain background information in the source image.

### 4.2.3 | MC-GAN

When synthesizing images, most text-to-image synthesis methods consider each output image as one single unit to characterize its semantic relevance to the texts. This is likely problematic because most images naturally consist of two crucial components: foreground and background. Without properly separating these two components, it's hard to characterize the semantics of an image if the whole image is treated as a single unit without proper separation.

In order to enhance the semantic relevance of the images, a multiconditional GAN (MC-GAN) (Park, Yoo, & Kwak, 2018) is proposed to synthesize a target image by combining the background of a source image and a text-described foreground object which does not exist in the source image. A unique feature of MC-GAN is that it proposes a synthesis block in which the background feature is extracted from the given image without nonlinear function (i.e., only using convolution and batch normalization) and the foreground feature is the feature map resulting the previous layer.

Because MC-GAN is able to properly model the background and foreground of the generated images, a unique strength of MC-GAN is that users are able to provide a base image and MC-GAN is able to preserve the background information of the base image to generate new images.

## 4.3 | Resolution enhancement GANs

Due to the fact that training GANs will be much difficult when generating high-resolution images, a two stage GAN (i.e., stackGAN) is proposed in which rough images (i.e., low-resolution images) are generated in stage-I and refined in stage-II. To further improve the quality of generated images, the second version of StackGAN (i.e., Stack++) is proposed to use multistage GANs to generate multiscale images. A color-consistency regularization term is also added into the loss to keep the consistency of images in different scales.

While stackGAN and StackGAN++ are both built on the global sentence vector, AttnGAN is proposed to use attention mechanism (i.e., deep attentional multimodal similarity model [DAMSM]) to model the multilevel information (i.e., word level and sentence level) into GANs. In the following, StackGAN, StackGAN++, and AttnGAN will be explained in detail.

Recently, dynamic memory generative adversarial network (DM-GAN) (Zhu, Pan, Chen, & Yang, 2019) which uses a dynamic memory component is proposed to focus on refining the initial generated image which is the key to the success of generating high-quality images.

### 4.3.1 | StackGAN

In 2017, Zhang et al. proposed a model for generating photorealistic images from text descriptions called StackGAN (stacked generative adversarial network) (Zhang, Xu, et al., 2017). In their work, they define a two-stage model that uses two cascaded GANs, each corresponding to one of the stages. The stage I GAN takes a text description as input, converts the text description to a text embedding containing several conditioning variables, and generates a low-quality 64x64 image with rough shapes and colors based on the computed conditioning variables. The stage II GAN then takes this low-quality stage I image as well as the same text embedding and uses the conditioning variables to correct and add more detail to the stage I result. The output of stage II is a photorealistic  $256 \times 256$  image that resembles the text description with compelling accuracy.

One major contribution of StackGAN is the use of cascaded GANs for text-to-image synthesis through a sketch-refinement process. By conditioning the stage II GAN on the image produced by the stage I GAN and text description,

the stage II GAN is able to correct defects in the stage I output, resulting in high-quality 256x256 images. Prior works have utilized “stacked” GANs to separate the image generation process into structure and style (Wang & Gupta, 2016), multiple stages each generating lower-level representations from higher-level representations of the previous stage (Huang et al., 2017), and multiple stages combined with a Laplacian pyramid approach (Denton, Chintala, Szlam, & Fergus, 2015), which was introduced for image compression by P. Burt and E. Adelson in 1983 and uses the differences between consecutive down samples of an original image to reconstruct the original image from its down sampled version (Burt & Adelson, 1983). However, these works did not use text descriptions to condition their generator models.

Conditioning augmentation is the other major contribution of StackGAN. Prior works transformed the natural language text description into a fixed text embedding containing static conditioning variables which were fed to the generator (Reed, Akata, Yan, et al., 2016). StackGAN does this and then creates a Gaussian distribution from the text embedding and randomly selects variables from the Gaussian distribution to add to the set of conditioning variables during training. This encourages robustness by introducing small variations to the original text embedding for a particular training image while keeping the training image that the generated output is compared to the same. The result is that the trained model produces more diverse images in the same distribution when using conditioning augmentation than the same model using a fixed text embedding (Zhang, Xu, et al., 2017).

### 4.3.2 | StackGAN++

Proposed by the same users as StackGAN, StackGAN++ is also a stacked GAN model, but organizes the generators and discriminators in a “tree-like” structure (Zhang, Xu, et al., 2018) with multiple stages. The first stage combines a noise vector and conditioning variables (with conditional augmentation introduced in Zhang, Xu, et al. (2017)) for input to the first generator, which generates a low-resolution image, 64x64 by default (this can be changed depending on the desired number of stages). Each following stage uses the result from the previous stage and the conditioning variables to produce gradually higher resolution images. These stages do not use the noise vector again, as the creators assume that the randomness it introduces is already preserved in the output of the first stage. The final stage produces a 256x256 high-quality image.

StackGAN++ introduces the joint conditional and unconditional approximation in their designs (Zhang, Xu, et al., 2018). The discriminators are trained to calculate the loss between the image produced by the generator and the conditioning variables (measuring how accurately the image represents the description) as well as the loss between the image and real images (probability of the image being real or fake). The generators then aim to minimize the sum of these losses, improving the final result.

### 4.3.3 | AttnGAN

Attentional generative adversarial network (AttnGAN) (Xu et al., 2018) is very similar, in terms of its structure, to StackGAN++ (Zhang, Xu, et al., 2018), discussed in the previous section, but some novel components are added. Like previous works (Reed et al., 2016; Zhang, Xu, et al., 2017, 2018), a text encoder generates a text embedding with conditioning variables based on the overall sentence. Additionally, the text encoder generates a separate text embedding with conditioning variables based on individual words. This process is optimized to produce meaningful variables using a bidirectional recurrent neural network (BRNN), more specifically bidirectional long short-term memory (LSTM) (Schuster & Paliwal, 2018), which, for each word in the description, generates conditions based on the previous word as well as the next word (bidirectional). The first stage of AttnGAN generates a low-resolution image based on the sentence-level text embedding and random noise vector. The output is fed along with the word-level text embedding to an “attention model,” which matches the word-level conditioning variables to regions of the stage I image, producing a word-context matrix. This is then fed to the next stage of the model along with the raw previous stage output. Each consecutive stage works in the same manner, but produces gradually higher resolution images conditioned on the previous stage.

Two major contributions were introduced in AttnGAN: the attentional generative network and the DAMSM (Zhang, Xu, et al., 2018). The attentional generative network matches specific regions of each stage's output image to conditioning variables from the word-level text embedding. This is a very worthy contribution, allowing each consecutive stage to focus on specific regions of the image independently, adding “attentional” details region by region as

opposed to the whole image. The DAMSM is also a key feature introduced by AttnGAN, which is used after the result of the final stage to calculate the similarity between the generated image and the text embedding at both the sentence level and the more fine-grained word level. Table 2 shows scores from different metrics for StackGAN, StackGAN++, AttnGAN, and HDGAN on the CUB, Oxford, and COCO datasets. The table shows that AttnGAN outperforms the other models in terms of IS on the CUB dataset by a small amount and greatly outperforms them on the COCO dataset.

**TABLE 1** A summary of different generative adversarial networks (GANs) and datasets used for validation

Method Names	Evaluation datasets				
	MNIST	Oxford-102	COCO	CUB	CIFAR-10
cGAN (Mirza & Osindero, 2014)	✓				
AC-GAN (Odena et al., 2017)					✓
TAC-GAN (Dash, Gamboa, Ahmed, Afzal, & Liwicki, 2017)		✓			
Text-SeGAN (Cha, Gown, & Kung, 2019)		✓			
GAN-INT-CLS (Reed, Akata, Yan, et al., 2016)		✓	✓	✓	
StackGAN (Zhang, Xu, et al., 2017)		✓	✓	✓	
StackGAN++ (Zhang, Xu, et al., 2017)		✓	✓	✓	
AttnGAN (Xu et al., 2018)		✓	✓	✓	
DC-GAN (Reed, Akata, Yan, et al., 2016)		✓	✓	✓	
HDGAN (Zhang, Xie, & Yang, 2018)		✓	✓	✓	
MirrorGAN (Qiao et al., 2019)			✓	✓	

Note: A ✓ symbol indicates that the model was evaluated using the corresponding dataset.

**TABLE 2** A summary of performance of different methods with respect to the three benchmark datasets and four performance metrics: Inception Score (IS), Frechet Inception Distance (FID), Human Classifier (HC), and SSIM scores

Methods	Datasets and Metrics											
	CUB				COCO				Oxford			
	IS	FID	HC	SSIM	IS	FID	HC	SSIM	IS	FID	HC	SSIM
DCGAN	2.88	68.79	2.76	—	7.88	60.82	1.82	—	2.66	79.55	1.84	—
GAN-INT-CLS	2.32	68.79	2.75	—	7.95	60.62	1.93	—	2.69	79.55	1.90	—
Dong-GAN	—	—	—	—	—	—	—	—	4.14	—	—	—
Paired-D-GAN	—	—	—	—	—	—	—	—	4.49	—	—	—
StackGAN	3.74	51.89	1.27	0.234	8.60	74.05	1.14	—	3.21	55.28	1.16	—
StackGAN++	4.09	15.30	1.17	—	8.40	81.59	1.55	—	3.27	48.68	1.27	—
AttnGAN	4.39	23.98	—	—	26.36	35.49	—	—	—	—	—	—
Obj-GAN	—	—	—	—	30.11	20.75	—	—	—	—	—	—
HDGAN	4.20	—	—	0.215	12.04	—	—	—	3.52	—	—	—
DM-GAN	4.82	16.09	—	—	31.06	32.64	—	—	—	—	—	—
TAC-GAN	—	—	—	—	—	—	—	—	3.50	—	—	—
Text-SeGAN	—	—	—	—	—	—	—	—	4.1	—	—	—
Scene Graph GAN	—	—	—	—	7.40	—	—	—	—	—	—	—
MirrorGAN	4.61	—	—	—	26.88	—	—	—	—	—	—	—

Note: The generative adversarial networks include DCGAN, GAN-INT-CLS, Dong-GAN, Paired-D-GAN, StackGAN, StackGAN++, AttnGAN, ObjGAN, HDGAN, DM-GAN, TAC-GAN, Text-SeGAN, Scene Graph GAN, and MirrorGAN. The three benchmark datasets include CUB, Oxford, and COCO datasets. A dash indicates that no data was found.

#### 4.3.4 | HDGAN

Hierarchically nested adversarial network (HDGAN) is a method proposed by Zhang, Xie, and Yang (2018), and its main objective is to tackle the difficult problem of dealing with photographic images from semantic text descriptions. These semantic text descriptions are applied on images from diverse datasets. This method introduces adversarial objectives nested inside hierarchically oriented networks (Zhang, Xie, & Yang, 2018). Hierarchical networks help regularize mid-level manifestations. In addition to regularize mid-level manifestations, it assists the training of the generator in order to capture highly complex still media elements. These elements are captured in statistical order to train the generator based on settings extracted directly from the image. The latter is an ideal scenario. However, this paper aims to incorporate a single-stream architecture. This single-stream architecture functions as the generator that will form an optimum adaptability toward the jointed discriminators. Once jointed discriminators are setup in an optimum manner, the single-stream architecture will then advance generated images to achieve a much higher resolution (Zhang, Xie, & Yang, 2018).

The main contributions of the HDGANs include the introduction of a visual-semantic similarity measure (Zhang, Xie, & Yang, 2018). This feature will aid in the evaluation of the consistency of generated images. In addition to checking the consistency of generated images, one of the key objectives of this step is to test the logical consistency of the end product (Zhang, Xie, & Yang, 2018). The end product in this case would be images that are semantically mapped from text-based natural language descriptions to each area on the picture, for example, a wing on a bird or petal on a flower. DL has created a multitude of opportunities and challenges for researchers in the computer vision AI field. Coupled with GAN and multimodal learning architectures, this field has seen tremendous growth (Li, Su, & Liu, 2018; Liu et al., 2018; Mirza & Osindero, 2015; Reed, Akata, Yan, et al., 2016; Wang & Gupta, 2016; Yang et al., 2017; Zhang, Tu, & Cui, 2017). Based on these advancements, HDGANs attempt to further extend some desirable and less common features when generating images from textual natural language (Zhang, Xie, & Yang, 2018). In other words, it takes sentences and treats them as a hierarchical structure. This has some positive and negative implications in most cases. For starters, it makes it more complex to generate compelling images. However, one of the key benefits of this elaborate process is the realism obtained once all processes are completed. In addition, one common feature added to this process is the ability to identify parts of sentences with bounding boxes. If a sentence includes common characteristics of a bird, it will surround the attributes of such bird with bounding boxes. In practice, this should happen if the desired image have other elements such as human faces (e.g., eyes, hair, etc.), flowers (e.g., petal size, color, etc.), or any other inanimate object (e.g., a table, a mug, etc.). Finally, HDGANs evaluated some of its claims on common ideal text-to-image datasets such as CUB, COCO, and Oxford-102 (Li, Su, & Liu, 2018; Liu et al., 2018; Mirza & Osindero, 2015; Reed, Akata, Yan, et al., 2016; Wang & Gupta, 2016; Yang et al., 2019; Zhang, Tu, & Cui, 2017; Zhang, Xie, & Yang, 2018). These datasets were first utilized on earlier works (Reed, Akata, Yan, et al., 2016), and most of them sport modified features such image annotations, labels, or descriptions. The qualitative and quantitative results reported by researchers in this study were far superior of earlier works in this same field of computer vision AI.

### 4.4 | Diversity enhancement GANs

In this subsection, we introduce text-to-image synthesis methods which try to maximize the diversity of the output images, based on the text descriptions.

#### 4.4.1 | Auxiliary classifier GANs (AC-GAN)

Two issues arise in the traditional GANs (Mirza & Osindero, 2015) for image synthesis: (1) Scalability problem: traditional GANs cannot predict a large number of image categories; and (2) diversity problem: images are often subject to one-to-many mapping, so one image could be labeled as different tags or being described using different texts. To address these problems, GAN conditioned on additional information, for example, cGAN, is an alternative solution. However, although cGAN and many previously introduced approaches are able to generate images with respect to the text descriptions, they often output images with similar types and visual appearance.

Slightly different from the cGAN, auxiliary classifier GANs (AC-GAN) (Odena et al., 2017) proposes to improve the diversity of output images by using an auxiliary classifier to control output images. The overall structure of AC-GAN is



shown in Figure 7c. In AC-GAN, every generated image is associated with a class label, in addition to the true/fake label which is commonly used in GAN or cGAN. The discriminator of AC-GAN not only outputs a probability distribution over sources (i.e., whether the image is true or fake), it also outputs a probability distribution over the class label (i.e., predict which class the image belong to).

By using an auxiliary classifier layer to predict the class of the image, AC-GAN is able to use the predicted class labels of the images to ensure that the output consists of images from different classes, resulting in diversified synthesis images. The results show that AC-GAN can generate images with high diversity.

#### 4.4.2 | TAC-GAN

Building on the AC-GAN, TAC-GAN (Dash, Gamboa, Ahmed, Liwicki, & Afzal, 2017) is proposed to replace the class information with textual descriptions as the input to perform the task of text to image synthesis. The architecture of TAC-GAN is shown in Figure 7d, which is similar to AC-GAN. Overall, the major difference between TAC-GAN and AC-GAN is that TAC-GAN conditions the generated images on text descriptions instead of on a class label. This design makes TAC-GAN more generic for image synthesis.

For TAC-GAN, it imposes restrictions on generated images in both texts and class labels. The input vector of TAC-GAN's generative network is built based on a noise vector and embedded vector representation of textual descriptions. The discriminator of TAC-GAN is similar to that of the AC-GAN, which not only predicts whether the image is fake or not, but also predicts the label of the images. A minor difference of TAC-GAN's discriminator, compared to that of the AC-GAN, is that it also receives text information as input before performing its classification.

The experiments and validations, on the Oxford-102 flowers dataset, show that the results produced by TAC-GAN are "slightly better" than other approaches, including GAN-INT-CLS and StackGAN.

#### 4.4.3 | Text-SeGAN

In order to improve the diversity of the output images, both AC-GAN and TAC-GAN's discriminators predict class labels of the synthesized images. This process likely enforces the semantic diversity of the images, but class labels are inherently restrictive in describing image semantics, and images described by text can be matched to multiple labels. Therefore, instead of predicting images' class labels, an alternative solution is to directly quantify their semantic relevance.

The architecture of Text-SeGAN is shown in Figure 7e. In order to directly quantify semantic relevance, Text-SeGAN (Cha, Gown, & Kung, 2019) adds a regression layer to estimate the semantic relevance between the image and text instead of a classifier layer of predicting labels. The estimated semantic relevance is a fractional value ranging between 0 and 1, with a higher value reflecting better semantic relevance between the image and text. Due to this unique design, an inherent advantage of Text-SeGAN is that the generated images are not limited to certain classes and are semantically matching to the text input.

Experiments and validations, on Oxford-102 flower dataset, show that Text-SeGAN can generate diverse images that are semantically relevant to the input text. In addition, the results of Text-SeGAN show improved IS compared to other approaches, including GAN-INT-CLS, StackGAN, TAC-GAN, and HDGAN.

#### 4.4.4 | MirrorGAN and Scene Graph GAN

Due to the inherent complexity of the visual images, and the diversity of text descriptions (i.e., same words could imply different meanings), it is difficult to precisely match the texts to the visual images at the semantic levels. For most methods we have discussed so far, they employ a direct text to image generation process, but there is no validation about how generated images comply with the text in a reverse fashion.

To ensure the semantic consistency and diversity, MirrorGAN (Qiao, Zhang, Xu, & Tao, 2019) employs a mirror structure, which reversely learns from generated images to output texts (an image-to-text process) to further validate whether generated are indeed consistent to the input texts. MirrorGAN includes three modules: a semantic text embedding module (STEM), a global-local collaborative attentive module for cascaded image generation (GLAM), and

a semantic text regeneration and alignment module (STREAM). The back to back Text-to-Image (T2I) and Image-to-Text (I2T) are combined to progressively enhance the diversity and semantic consistency of the generated images.

In order to enhance the diversity of the output image, Scene Graph GAN (Johnson, Gupta, & Fei-Fei, 2018) proposes to use visual scene graphs to describe the layout of the objects, allowing users to precisely specify the relationships between objects in the images. In order to convert the visual scene graph as input for GAN to generate images, this method uses graph convolution to process input graphs. It computes a scene layout by predicting bounding boxes and segmentation masks for objects. After that, it converts the computed layout to an image with a cascaded refinement network.

## 4.5 | Motion enhancement GANs

Instead of focusing on generating static images, another line of text-to-image synthesis research focuses on generating videos (i.e., sequences of images) from texts. In this context, the synthesized videos are often useful resources for automated assistance or story telling.

### 4.5.1 | ObamaNet and T2S

One early/interesting work of motion enhancement GANs is to generate spoofed speech and lip-sync videos (or talking face) of Barack Obama (i.e., ObamaNet) based on text input (Kumar, Sotelo, Kumar, de Brébisson, & Bengio, 2017). This framework is consisted of three parts, that is, text to speech using “Char2Wav,” mouth shape representation synced to the audio using a time-delayed LSTM and “video generation” conditioned on the mouth shape using “U-Net” architecture. Although the results seem promising, ObamaNet only models the mouth region and the videos are not generated from noise which can be regarded as video prediction other than video generation.

Another meaningful trial of using synthesized videos for automated assistance is to translate spoken language (e.g., text) into sign language video sequences (i.e., T2S) (Stoll, Camgöz, Hadfield, & Bowden, 2018). This is often achieved through a two-step process: converting texts as meaningful units to generate images, followed by a learning component to arrange images into sequential order for best representation. More specifically, using RNN-based machine translation methods, texts are translated into sign language gloss sequences. Then, glosses are mapped to skeletal pose sequences using a lookup-table. To generate videos, a conditional DCGAN with the input of concatenation of latent representation of the image for a base pose and skeletal pose information is built.

### 4.5.2 | Text-to-video model (T2V)

In Li, Min, Shen, Carlson, and Carin (2018), a text-to-video model (*T2V*) is proposed based on the *cGAN* in which the input is the isometric Gaussian noise with the text-gist vector served as the generator. A key component of generating videos from text is to train a conditional generative model to extract both static and dynamic information from text, followed by a hybrid framework combining a variational auto encoder (VAE) and a GAN.

More specifically, *T2V* relies on two types of features, static features and dynamic features, to generate videos. Static features, called “gist” are used to sketch text-conditioned background color and object layout structure. Dynamic features, on the other hand, are considered by transforming input text into an image filter which eventually forms the video generator which consists of three entangled neural networks. The text-gist vector is generated by a gist generator which maintains static information (e.g., background) and a *text2filter* which captures the dynamic information (i.e., actions) in the text to generate videos.

As demonstrated in the paper (Li, Min, et al., 2018), the generated videos are semantically related to the texts, but have a rather low quality (e.g., only  $64 \times 64$  resolution).

### 4.5.3 | Storyland

Different from *T2V* which generates videos from a single text, *StoryGAN* aims to produce dynamic scenes consistent of specified texts (i.e., story written in a multi-sentence paragraph) using a sequential GAN model (Li, Gan, et al., 2019).

Story encoder, context encoder, and discriminators are the main components of this model. By using stochastic sampling, the story encoder intends to learn a low-dimensional embedding vector for the whole story to keep the continuity of the story. The context encoder is proposed to capture contextual information during sequential image generation based on a deep RNN. Two discriminators of StoryGAN are image discriminator which evaluates the generated images and story discriminator which ensures the global consistency.

The experiments and comparisons, on CLEVR dataset and Pororo cartoon dataset which are originally used for visual question answering, show that StoryGAN improves the generated video quality in terms of structural similarity index (SSIM), visual quality, consistence, and relevance (the last three measure are based on human evaluation).

## 5 | GAN-BASED TEXT-TO-IMAGE SYNTHESIS APPLICATIONS, BENCHMARK, AND EVALUATION AND COMPARISONS


### 5.1 | Text-to-image synthesis applications

Computer vision applications have strong potential for industries including but not limited to the medical, government, military, entertainment, and online social media fields (Asmuth et al., 1998; Fang, Wang, & Tang, 2018; Hong, Yang, Choi, & Lee, 2018; Mansimov, Parisotto, Ba, & Salakhutdinov, 2016; Nie, Trullo, Petitjean, Ruan, & Shen, 2016; Wu et al., 2017). Text-to-image synthesis is one such application in computer vision AI that has become the main focus in recent years due to its potential for providing beneficial properties and opportunities for a wide range of applicable areas.

Text-to-image synthesis is an application byproduct of deep convolutional decoder networks in combination with GANs (Reed, Akata, Yan, et al., 2016; Wu et al., 2017; Xu et al., 2018). Deep convolutional networks have contributed to several breakthroughs in image, video, speech, and audio processing. This learning method intends, among other possibilities, to help translate sequential text descriptions to images supplemented by one or many additional methods. Algorithms and methods developed in the computer vision field have allowed researchers in recent years to create realistic images from plain sentences. Advances in the computer vision, deep convolutional nets, and semantic units have shined light and redirected focus to this research area of text-to-image synthesis, having as its prime directive: to aid in the generation of compelling images with as much fidelity to text descriptions as possible.

To date, models for generating synthetic images from textual natural language in research laboratories at universities and private companies have yielded compelling images of flowers and birds (Reed, Akata, Yan, et al., 2016). Although flowers and birds are the most common objects studied thus far, research has been applied to other classes as well. For example, there have been studies focused solely on human faces (Reed, Akata, Yan, et al., 2016; Wang, Chang, Cheng, Jin, & Cheng, 2018; Wu et al., 2017; Yin, Fu, Sigaly, & Xue, 2017).

It's a fascinating time for computer vision AI and DL researchers and enthusiasts. The consistent advancement in hardware, software, and contemporaneous development of computer vision AI research disrupts multiple industries.



**FIGURE 10** Performance comparison between 14 generative adversarial networks (GANs) with respect to their inception scores (ISs)

These advances in technology allow for the extraction of several data types from a variety of sources. For example, image data captured from a variety of photo-ready devices, such as smart phones, and online social media services opened the door to the analysis of large amounts of media datasets (Fang et al., 2018). The availability of large media datasets allow new frameworks and algorithms to be proposed and tested on real-world data.

## 5.2 | Text-to-image synthesis benchmark datasets

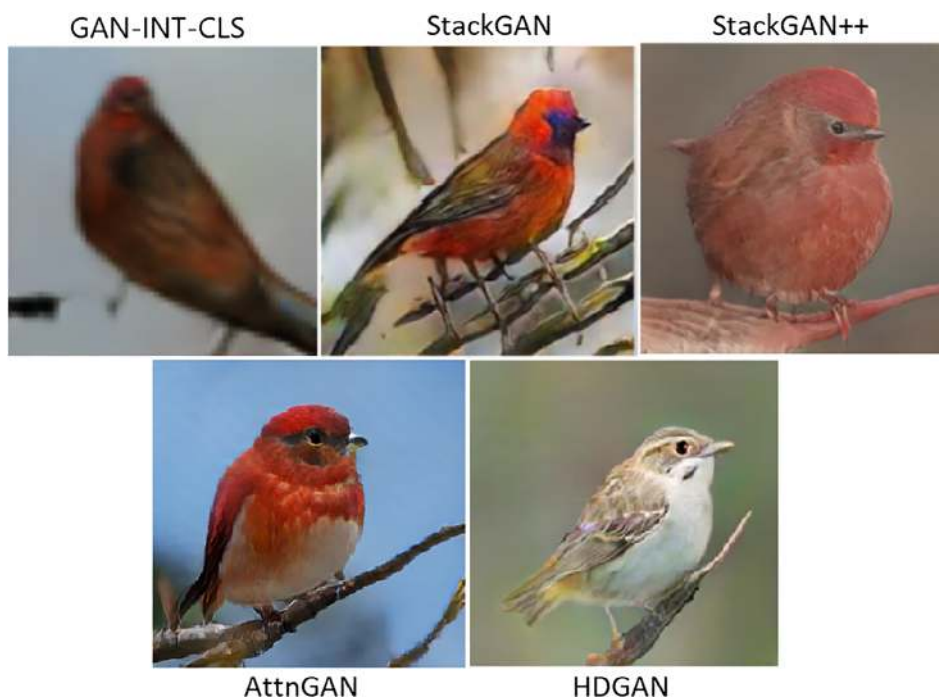
A summary of some reviewed methods and benchmark datasets used for validation is reported in Table 1. In addition, the performance of different GANs with respect to the benchmark datasets and performance metrics is reported in Table 2.

In order to synthesize images from text descriptions, many frameworks have taken a minimalistic approach by creating small and background-less images (Mao et al., 2017). In most cases, the experiments were conducted on simple datasets, initially containing images of birds and flowers. Reed, Akata, Yan, et al. (2016) contributed to these datasets by adding corresponding natural language text descriptions to subsets of the CUB, MSCOCO, and Oxford-102 datasets, which facilitated the work on text-to-image synthesis for several papers released more recently.

While most DL algorithms use MNIST (LeCun & Cortes, 2010) dataset as the benchmark, there are three main datasets that are commonly used for evaluation of proposed GAN models for text-to-image synthesis: CUB (Wang et al., 2011), Oxford (Nilsback & Zisserman, 2008), COCO (Lin et al., 2015), and CIFAR-10 (Krizhevsky, 2009). CUB (Wang et al., 2011) contains 200 birds with matching text descriptions and Oxford (Nilsback & Zisserman, 2008) contains 102 categories of flowers with 40–258 images each and matching text descriptions. These datasets contain individual objects, with the text description corresponding to that object, making them relatively simple. COCO (Lin et al., 2015) is much more complex, containing 328 k images with 91 different object types. CIFAR-10 (Krizhevsky, 2009) dataset consists of 60,000  $32 \times 32$  color images in 10 classes, with 6,000 images per class. In contrast to CUB and Oxford, whose images each contain an individual object, COCO's images may contain multiple objects, each with a label, so there are many labels per image. The total number of labels over the 328 k images is 2.5 million (Lin et al., 2015).

## 5.3 | Text-to-image synthesis benchmark evaluation metrics

Several evaluation metrics are used for judging the images produced by text-to-image GANs. Proposed by Salimans et al. (2016), ISs calculates the entropy (randomness) of the conditional distribution, obtained by applying the inception



**FIGURE 11** Examples of best images of “birds” generated by GAN-INT-CLS, StackGAN, StackGAN++, AttnGAN, and HDGAN.(Reprinted with permission from Zhang, Xu, et al. (2017, 2018); Xu et al. (2018), and Zhang, Xie, and Yang (2018), respectively. Copyright 2020 IEEE)

**FIGURE 12** Examples of best images of “a plate of vegetables” generated by GAN-INT-CLS, StackGAN, StackGAN++, AttnGAN, and HDGAN. (Reprinted with permission from Zhang, Xu, et al. (2017, 2018); Xu et al. (2018), and Zhang, Xie, and Yang (2018), respectively. Copyright 2020 IEEE)



model introduced in Szegedy, Vanhoucke, Ioffe, and Shlens (2016), and marginal distribution of a large set of generated images, which should be low and high, respectively, for meaningful images. Low entropy of conditional distribution means that the evaluator is confident that the images came from the data distribution, and high entropy of the marginal distribution means that the set of generated images is diverse, which are both desired features. The IS score is then computed as the KL-divergence between the two entropies. Fully Convolutional Network (FCN)-scores (Isola et al., 2018) are computed in a similar manner, relying on the intuition that realistic images generated by a GAN should be able to be classified correctly by a classifier trained on real images of the same distribution. Therefore, if the FCN classifier classifies a set of synthetic images accurately, the image is probably realistic, and the corresponding GAN gets a high FCN score. FID (Heusel, Ramsauer, Unterthiner, Nessler, & Hochreiter, 2018) is the other commonly used evaluation metric, and takes a different approach, actually comparing the generated images to real images in the distribution. A high FID means there is little relationship between statistics of the synthetic and real images and vice versa, so lower FIDs are better.

The performance of different GANs with respect to the benchmark datasets and performance metrics is reported in Table 2. In addition, Figure 10 further lists the performance of 14 GANs with respect to their ISs.

## 5.4 | GAN-based text-to-image synthesis results comparison

While we gathered all the data we could find on scores for each model on the CUB, Oxford, and COCO datasets using IS, FID, FCN, and human classifiers, we unfortunately were unable to find certain data for AttnGAN and HDGAN (missing in Table 2). The best evaluation we can give for those with missing data is our own opinions by looking at examples of generated images provided in their papers. In this regard, we observed that HDGAN produced relatively better visual results on the CUB and Oxford datasets while AttnGAN produced far more impressive results than the rest on the more complex COCO dataset. This is evidence that the attentional model and DAMSM introduced by AttnGAN are very effective in producing high-quality images. Examples of the best results of birds and plates of vegetables generated by each model are presented in Figures 11 and 12, respectively.

In terms of IS, which is the metric that was applied to majority models except DC-GAN, the results in Table 2 show that StackGAN++ only showed slight improvement over its predecessor, StackGAN, for text-to-image synthesis. However, StackGAN++ did introduce a very worthy enhancement for unconditional image generation by organizing the generators and discriminators in a “tree-like” structure. This indicates that revising the structures of the discriminators and/or generators can bring a moderate level of improvement in text-to-image synthesis.

In addition, the results in Table 2 also show that DM-GAN (Zhu et al., 2019) has the best performance, followed by Obj-GAN (Li, Gan, et al., 2019). Note that both DM-GAN and Obj-GAN are most recently developed methods in the field (both published in 2019), indicating that research in text to image synthesis is continuously improving the results for better visual perception and interception. Technical wise, DM-GAN (Zhu et al., 2019) is a model using dynamic memory to refine fuzzy image contents initially generated from the GAN networks. A memory writing gate is used for DM-GAN to select important text information and generate images based on he selected text accordingly. On the other hand, Obj-GAN (Li, Zhang, et al., 2019) focuses on object centered text-to-image synthesis. The proposed framework of Obj-GAN consists of a layout generation, including a bounding box generator and a shape generator, and an object-driven attentive image generator. The designs and advancement of DM-GAN and Obj-GAN indicate that research in text-to-image synthesis is advancing to put more emphasis on the image details and text semantics for better understanding and perception.

## 5.5 | Notable mentions

It is worth noting that although this survey mainly focuses on text-to-image synthesis, there have been other applications of GANs in broader image synthesis field that we found fascinating and worth dedicating a small section to. For example, Yin et al. (2017) used Sem-Latent GANs to generate images of faces based on facial attributes, producing impressive results that, at a glance, could be mistaken for real faces. Xu et al. (2018); Fang et al. (2018), and Karpathy and Fei-Fei (2018) demonstrated great success in generating text descriptions from images (image captioning) with great accuracy, with Xu et al. (2018) using an attention-based model that automatically learns to focus on salient objects and Karpathy and Fei-Fei (2016) using deep visual-semantic alignments. Finally, there is a contribution made by StackGAN++ that was not mentioned in the dedicated section due to its relation to unconditional image generation as opposed to conditional, namely a color-regularization term (Zhang, Xu, et al., 2018). This additional term aims to keep the samples generated from the same input at different stages more consistent in color, which resulted in significantly better results for the unconditional model.

## 6 | CONCLUSION

The recent advancement in text-to-image synthesis research opens the door to several compelling methods and architectures. The main objective of text-to-image synthesis initially was to create images from simple labels, and this objective later scaled to natural languages. In this paper, we reviewed novel methods that generate, in our opinion, the most visually rich and photorealistic images, from text-based natural language. These generated images often rely on GANs, deep convolutional decoder networks, and multimodal learning methods.

In the paper, we first proposed a taxonomy to organize GAN-based text-to-image synthesis frameworks into four major groups: semantic enhancement GANs, resolution enhancement GANs, diversity enhancement GANs, and motion enhancement GANs. The taxonomy provides a roadmap to show the motivations, architectures, and difference of different methods, and also outlines their evolution timeline and relationships. Following the proposed taxonomy, we reviewed important features of each method and their architectures. We indicated the model definition and key contributions from some advanced GAN frameworks, including StackGAN, StackGAN++, AttnGAN, DC-GAN, AC-GAN, TAC-GAN, HDGAN, Text-SeGAN, StoryGAN and so on. Many of the solutions surveyed in this paper tackled the highly complex challenge of generating photorealistic images beyond swatch size samples. In other words, beyond the work of (Reed, Akata, Yan, et al., 2016) in which images were generated from text in 64×64 tiny swatches. Last, all methods were evaluated on datasets that included birds, flowers, humans, and other miscellaneous elements. We were also able to allocate some important papers that were as impressive as the papers we finally surveyed. Though, these notable papers have yet to contribute directly or indirectly to the expansion of the vast computer vision AI field. Looking into the future, an excellent extension from the works surveyed in this paper would be to give more independence to the several learning methods (e.g., less human intervention) involved in the studies as well as increasing the size of the output images.

## ACKNOWLEDGMENTS

This work is sponsored in part by the US National Science Foundation (NSF) through Grants No. IIS-1763452 and CNS-1828181, and by National Natural Science Foundation of China (NSFC) under Grant No. 91646204.

## CONFLICT OF INTEREST

The authors declare that there is no conflict of interest regarding the publication of this article.

## AUTHOR CONTRIBUTIONS

**Jorge Agnese:** Methodology; writing-original draft, review, and editing. **Jonathan Herrera:** Methodology; software; writing-original draft, review, and editing. **Haicheng Tao:** Methodology; writing-original draft, review, and editing. **Xingquan Zhu:** Funding acquisition; project administration; supervision; writing-review and editing.

## ORCID

Xingquan Zhu  <https://orcid.org/0000-0003-4129-9611>

## RELATED WIREs ARTICLES

[Deep learning for sentiment analysis: A survey](#)

## FURTHER READING

- Li, Y., Gan, Z., Shen, Y., Liu, J., Cheng, Y., Wu, Y., ... Gao, J. (2019). Storygan: A sequential conditional gan for story visualization. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 6329–6338). Long Beach, CA: IEEE Computer Society.
- Schuster, M., & Paliwal, K. (1997). Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11), 2673–2681.
- 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.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhutdinov, R., ... Bengio, Y. (2015). Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning* (pp. 2048–2057). Lille, France. Retrieved from JMLR.org
- 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 *2018 IEEE/CVF conference on computer vision and pattern recognition* (pp. 1316–1324). Salt Lake City, UT: IEEE Computer Society.
- Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., & Metaxas, D. (2019). StackGAN++: Realistic image synthesis with stacked generative adversarial networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41, 1947–1962.

## REFERENCES

- Asmuth, J., Dixon, D., Hanna, K., Hsu, S. C., Kumar, R., Paragano, V., ... Sawhney, H. (1998). Multimedia applications of computer vision. In *Proceedings fourth IEEE workshop on applications of computer vision* (pp. 290–291). Princeton, NJ: IEEE Computer Society.
- Baltrusaitis, T., Ahuja, C., & Morency, L. (2017). Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41, 423–443.
- Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798–1828.
- Burt, P., & Adelson, E. (1983). The laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, 31(4), 532–540.
- Cha, M., Gwon, Y. L., & Kung, H. T. (2019). Adversarial learning of semantic relevance in text to image synthesis. In *The thirty-third AAAI conference on artificial intelligence* (pp. 3272–3279). Long Beach, CA: IEEE Computer Society.
- Cha, M., Gwon, Y., & Kung, H. (2017). Adversarial nets with perceptual losses for text-to-image synthesis. In *2017 IEEE 27th international workshop on machine learning for signal processing (MLSP)* (pp. 1–6). Tokyo, Japan: IEEE.
- Cha, M., Gwon, Y. L., & Kung, H. (2019). Adversarial learning of semantic relevance in text to image synthesis. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 33, pp. 3272–3279). Honolulu, HI: AAAI Press.
- Chen, J., Shen, Y., Gao, J., Liu, J., & Liu, X. (2018). Language-based image editing with recurrent attentive models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 8721–8729). Salt Lake City, UT: IEEE Computer Society.
- Chen, Y., Lai, Y.-K., & Liu, Y.-J. (2018). Cartoongan: Generative adversarial networks for photo cartoonization. In *IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 9465–9474). Salt Lake City, UT: IEEE Computer Society.
- Creswell, A., White, T., Dumoulin, V., Arulkumaran, K., Sengupta, B., & Bharath, A. (2018). Generative adversarial networks: An overview. In *IEEE signal processing magazine* (pp. 53–65). IEEE.
- Dash, A., Gamboa, J. C. B., Ahmed, S., Afzal, M. Z., & Liwicki, M. (2017). Tac-gan-text conditioned auxiliary classifier generative adversarial network. *CoRR, arXiv*, 1703.06412.
- Dash, A., Gamboa, J. C. B., Ahmed, S., Liwicki, M., & Afzal, M. Z. (2017). Tac-gan-text conditioned auxiliary classifier generative adversarial network. *arXiv preprint arXiv:1703.06412*.
- Denton, E. L., Chintala, S., Szlam, A., & Fergus, R. (2015). Deep generative image models using a Laplacian pyramid of adversarial networks. In *Advances in neural information processing systems, Montreal, Quebec, Canada* (pp. 1486–1494). Red Hook, NY: Curran Associates, Inc.

- Dong, H., Yu, S., Wu, C., & Guo, Y. (2017). Semantic image synthesis via adversarial learning. In *Proceedings of the IEEE international conference on computer vision* (pp. 5706–5714). Honolulu, HI: IEEE Computer Society.
- Dong, H., Zhang, J., McIlwraith, D., & Guo, Y. (2017). I2t2i: Learning text to image synthesis with textual data augmentation. In *IEEE international conference on image processing (ICIP)* (pp. 2015–2019). Beijing, China: IEEE.
- Durugkar, I., Gemp, I., & Mahadevan, S. (2017). Generative multi-adversarial networks. In *Proceedings of the international conference on learning representations*. Toulon, France. Retrieved from OpenReview.net
- Elgammal, A., Liu, B., Elhoseiny, M., & Mazzone, M. (2017). Can: Creative adversarial networks generating “art” by learning about styles and deviating from style norms. *CoRR, arXiv, 1706.07068v1*.
- Fang, F., Wang, H., & Tang, P. (2018). Image captioning with word level attention. In *25th IEEE international conference on image processing (ICIP)* (pp. 1278–1282). Athens, Greece: IEEE.
- Gao, L., Chen, D., Song, J., Xu, X., Zhang, D., & Shen, H. T. (2019). Perceptual pyramid adversarial networks for text-to-image synthesis. In *Proceedings of the association for the advancement of artificial intelligence (AAAI)* (Vol. 33, pp. 8312–8319). Honolulu, HI: AAAI Press.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... Bengio, Y. (2014). Generative adversarial nets. In *Proceedings of the 27th international conference on neural information processing systems, Montreal, Canada* (pp. 2672–2680). Red Hook, NY: Curran Associates, Inc.
- Han, H., Li, Y., & Zhu, X. (2019). Convolutional neural network learning for generic data classification. *Information Sciences, 477*, 448–465.
- Haynes, M., Norton, A., McParland, A., & Cooper, R. (2018). Speech-to-text for broadcasters, from research to implementation. *SMPTE Motion Imaging Journal, 127*(2), 27–33.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., & Hochreiter, S. (2018). Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Proceedings of the 31st international conference on neural information processing systems, Long Beach, California, USA* (pp. 6629–6640). Red Hook, NY, USA: Curran Associates Inc.
- Hong, S., Yang, D., Choi, J., & Lee, H. (2018). Inferring semantic layout for hierarchical text-to-image synthesis. In *2018 IEEE/CVF conference on computer vision and pattern recognition* (pp. 7986–7994). Salt Lake City, UT: IEEE Computer Society.
- Hong, Y., Hwang, U., Yoo, J., & Yoon, S. (2019). How generative adversarial networks and their variants work: An overview. *ACM Computing Surveys (CSUR), 52*(1), 10:1–10:43.
- Huang, H., Yu, P., & Wang, C. (2018). An introduction to image synthesis with generative adversarial nets. *CoRR, arXiv, 1803.04469v2*.
- Huang, X., Li, Y., Poursaeed, O., Hopcroft, J., & Belongie, S. (2017). Stacked generative adversarial networks. In *IEEE conference on computer vision and pattern recognition* (pp. 1866–1875). Honolulu, HI: IEEE Computer Society.
- Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. (2018). Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1125–1134). Honolulu, HI: IEEE Computer Society.
- 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). Amsterdam, The Netherlands: Springer.
- Johnson, J., Gupta, A., & Fei-Fei, L. (2018). Image generation from scene graphs. In *2018 IEEE conference on computer vision and pattern recognition* (pp. 1219–1228). Long Beach, CA: IEEE Computer Society.
- Karpathy, A., & Fei-Fei, L. (2016). Deep visual-semantic alignments for generating image descriptions. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 39*(4), 664–676.
- Krizhevsky, A. (2009). *Learning multiple layers of features from tiny images*. Technical report, Citeseer.
- Kumar, R., Sotelo, J., Kumar, K., de Brébisson, A., & Bengio, Y. (2017). Obamanet: Photo-realistic lip-sync from text. *arXiv preprint arXiv: 1801.01442*.
- LeCun, Y., & Cortes, C. (2010). *MNIST handwritten digit database*. Retrieved from <http://yann.lecun.com/exdb/mnist/>
- Li, C., Su, Y., & Liu, W. (2018). Text-to-text generative adversarial networks. In *International joint conference on neural networks (IJCNN)* (pp. 1–7). Rio de Janeiro, Brazil: IEEE.
- Li, C., Wang, Z., & Qi, H. (2018). Fast-converging conditional generative adversarial networks for image synthesis. In *2018 25th IEEE international conference on image processing (ICIP)* (pp. 2132–2136). Athens, Greece: IEEE.
- Li, W., Zhang, P., Zhang, L., Huang, Q., He, X., Lyu, S., & Gao, J. (2019). Object-driven text-to-image synthesis via adversarial training. In *The IEEE conference on computer vision and pattern recognition (CVPR)* (pp. 12174–12182). Long Beach, CA: IEEE Computer Society.
- Li, Y., Min, M. R., Shen, D., Carlson, D., & Carin, L. (2018). Video generation from text. In *Thirty-second AAAI conference on artificial intelligence* (pp. 7065–7072). New Orleans, LA.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., ... Zitnick, C. L. (2015). Microsoft coco: Common objects in context. In *European conference on computer vision* (pp. 740–755). Zurich, Switzerland: Springer.
- Liu, X., Meng, G., Xiang, S., & Pan, C. (2018). Semantic image synthesis via conditional cycle-generative adversarial networks. In *24th international conference on pattern recognition (ICPR)* (pp. 988–993). Beijing, China: IEEE Computer Society.
- Mansimov, E., Parisotto, E., Ba, J. L., & Salakhutdinov, R. (2016). Generating images from captions with attention. *CoRR, arXiv, 1502.03044v3*.
- Mao, X., Li, Q., Xie, H., Lau, R. Y. K., Wang, Z., & Smolley, S. P. (2017). Least squares generative adversarial networks. In *2017 IEEE international conference on computer vision (ICCV)* (pp. 2813–2821). Venice, Italy: IEEE Computer Society.
- Mirza, M., & Osindero, S. (2014). Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.
- Ngiam, J., Khosla, A., Kim, M., Nam, J., Lee, H., & Ng, A. Y. (2011). Multimodal deep learning. In *Proceedings of the 28th international conference on international conference on machine learning* (pp. 689–696). Bellevue, WA: Omnipress.



- Nguyen, T. D., Le, T., Vu, H., & Phung, D. (2017). Dual discriminator generative adversarial nets. In *Proceedings of the 31st international conference on neural information processing systems, Long Beach, California, USA* (pp. 2667–2677). Red Hook, NY: Curran Associates Inc.
- Nie, D., Trullo, R., Lian, J., Petitjean, C., Ruan, S., Wang, Q., & Shen, D. (2016). Medical image synthesis with context-aware generative adversarial networks. In *Medical image computing and computer assisted intervention* (pp. 417–425). Quebec City, QC, Canada: Springer International.
- Nilsback, M., & Zisserman, A. (2008). Automated flower classification over a large number of classes. In *Proceedings of the Indian conference on computer vision, graphics and image processing* (pp. 722–729). Bhubaneswar, India: IEEE Computer Society.
- Odena, A., Olah, C., & Shlens, J. (2017). Conditional image synthesis with auxiliary classifier gans. In *Proceedings of the 34th international conference on machine learning* (pp. 2642–2651). Sydney, NSW, Australia. Retrieved from JMLR.org
- Park, H., Yoo, Y., & Kwak, N. (2018). Mc-gan: Multi-conditional generative adversarial network for image synthesis. In *British machine vision conference* (p. 76). Newcastle, England: BMVA Press.
- Qiao, T., Zhang, J., Xu, D., & Tao, D. (2019). Mirrorgan: Learning text-to-image generation by redescription. In *The IEEE conference on computer vision and pattern recognition (CVPR)* (pp. 1505–1514). Long Beach, CA: IEEE Computer Society.
- Radford, A., Metz, L., & Chintala, S. (2016). Unsupervised representation learning with deep convolutional generative adversarial networks. *CoRR, arXiv*, **1511.06434v2**.
- Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., & Lee, H. (2016). Generative adversarial text to image synthesis. In *Proceedings of the 33rd international conference on machine learning* (pp. 1060–1069). New York, NY. Retrieved from JMLR.org
- Reed, S. E., Akata, Z., Mohan, S., Tenka, S., Schiele, B., & Lee, H. (2016). Learning what and where to draw. In *Advances in neural information processing systems, Barcelona, Spain* (pp. 217–225). Red Hook, NY: Curran Associates, Inc.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., & Chen, X. (2016). Improved techniques for training gans. In *Advances in neural information processing systems, Barcelona, Spain* (pp. 2234–2242). Red Hook, NY: Curran Associates, Inc.
- Stoll, S., Camgöz, N. C., Hadfield, S., & Bowden, R. (2018). Sign language production using neural machine translation and generative adversarial networks. In *British machine vision conference 2018* (pp. 304). Newcastle, England: BMVA Press.
- Szegedy, C., Vanhoucke, V., Ioffe, S., & Shlens, J. (2016). Rethinking the inception architecture for computer vision. In *IEEE conference on computer vision and pattern recognition* (pp. 2818–2826). Las Vegas, NV: IEEE Computer Society.
- Vo, D. M., & Sugimoto, A. (2018). Paired-d gan for semantic image synthesis. In *Asian conference on computer vision* (pp. 468–484). Perth, Australia: Springer.
- Wang, K., Gou, C., Duan, Y., Lin, Y., Zheng, X., & Wang, F. (2017). Generative adversarial networks: Introduction and outlook. *IEEE/CAA Journal of Automatica Sinica*, *4*(4), 588–598.
- Wang, K., & Wan, X. (2018). Sentigan: Generating sentimental texts via mixture adversarial networks. In *Proceedings of the twenty-seventh international joint conference on artificial intelligence* (pp. 4446–4452). Stockholm, Sweden. Retrieved from ijcai.org
- Wang, X., & Gupta, A. (2016). Generative image modeling using style and structure adversarial networks. In *Computer vision—ECCV 2016, Amsterdam, The Netherlands* (pp. 318–335). Cham, Switzerland: Springer International.
- Wang, Y., Chang, L., Cheng, Y., Jin, L., & Cheng, Z. (2018). Text2sketch: Learning face sketch from facial attribute text. In *IEEE international conference on image processing (ICIP)* (pp. 669–673). Athens, Greece: IEEE.
- Wu, X., Xu, K., & Hall, P. (2017). A survey of image synthesis and editing with generative adversarial networks. *Tsinghua Science and Technology*, *22*(6), 660–674.
- Yan, X., Yang, J., Sohn, K., & Lee, H. (2016). Attribute2image: Conditional image generation from visual attributes. In *Computer vision—ECCV 2016* (pp. 776–791). Amsterdam, The Netherlands: Springer International.
- Yan, Z., Zhang, H., Wang, B., Paris, S., & Yu, Y. (2016). Automatic photo adjustment using deep neural networks. *ACM Transactions on Graphics (TOG)*, *35*(2), 11:1–11:15.
- Yang, J., Kannan, A., Batra, D., & Parikh, D. (2017). Lr-gan: Layered recursive generative adversarial networks for image generation. In *5th international conference on learning representations*. Toulon, France. Retrieved from OpenReview.net
- Yang, M., Zhao, W., Xu, W., Feng, Y., Zhao, Z., Chen, X., & Lei, K. (2019). Multitask learning for cross-domain image captioning. *IEEE Transactions on Multimedia*, *21*, 1047–1061.
- Yin, W., Fu, Y., Sigaly, L., & Xue, X. (2017). Semi-latent gan: Learning to generate and modify facial images from attributes. *CoRR, arXiv*, **1703.01560v3**.
- Zhang, D., Yin, J., Zhu, X., & Zhang, C. (2018). Network representation learning: A survey. *IEEE Transactions on Big Data*. <https://doi.org/10.1109/TBDDATA.2018.2850013>
- Zhang, G., Tu, E., & Cui, D. (2017). Stable and improved generative adversarial nets (gans): A constructive survey. In *Proceedings of the international conference on image processing* (pp. 1871–1875). Beijing, China: IEEE.
- Zhang, H., Xu, T., Li, H., Zhang, S., Wang, Z., Huang, X., & Metaxas, D. (2017). Least squares generative adversarial networks. In *2017 IEEE international conference on computer vision (ICCV)* (pp. 2813–2821). Venice, Italy: IEEE Computer Society.
- Zhang, S., Zhai, J., Luo, D., Zhan, Y., & Chen, J. (2018). Recent advance on generative adversarial networks. In *Proceedings of the 2018 International conference on machine learning and cybernetics* (pp. 69–74). Chengdu, China: IEEE.
- Zhang, Z., Xie, Y., & Yang, L. (2018). Photographic text-to-image synthesis with a hierarchically-nested adversarial network, Salt Lake City, UT, USA. In *2018 IEEE/CVF conference on computer vision and pattern recognition* (pp. 6199–6208). IEEE Computer Society.
- 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). Long Beach, CA: IEEE Computer Society.

Zhu, X., Goldberg, A. B., Eldawy, M., Dyer, C. R., & Strock, B. (2007). A text-to-picture synthesis system for augmenting communication. In *Proceedings of the 22nd national conference on artificial intelligence* (pp. 1590–1595). Vancouver, British Columbia, Canada: AAAI Press.

**How to cite this article:** Agnese J, Herrera J, Tao H, Zhu X. A survey and taxonomy of adversarial neural networks for text-to-image synthesis. *WIREs Data Mining Knowl Discov.* 2020;10:e1345. <https://doi.org/10.1002/widm.1345>