

Image Synthesis and Editing with Generative Adversarial Networks (GANs): A Review.

Wanwan Li

Department of Computer Science

George Mason University

Fairfax, VA, US

wli17@gmu.edu

Abstract—Recently, as many deep learning models are emerging, deep learning has achieved great success in the field of artificial intelligence(AI). Especially, the Generative adversarial networks (GANs) based on zero-sum game theory has become a new research hot spot in the field of deep learning. The significance of the GAN model is that it can generate realistic data through unsupervised learning. Based on the conceptual and theoretical framework of the generative adversarial network, GANs models and their application result in tremendous success among different areas, especially in image synthesis and editing. This paper visualizes the data structures of various kinds of GANs models in 3D and discusses the variational GAN models with respect to their improvements in the applications. As the GANs have superior learning ability, strong plasticity, great potential for improvement, and a wide application range, this paper prospects the possible applications of the GANs in the near future.

Index Terms—Image Synthesis, Image Editing, Neural Networks, Deep Learning, Generative Adversarial Networks

I. INTRODUCTION

Generative adversarial networks (GANs) have become a hot research direction in the field of artificial intelligence(AI). The basic idea of GANs is derived from the two-person zero-sum game in game theory. It consists of two different neural networks: a generator and a discriminator. Typically, it is trained by means of adversarial unsupervised learning. The purpose of GAN is to estimate the potential distribution of data samples and generate new data samples. In the fields of image processing and visual computing, speech and natural language processing, information security, chess games, etc, GANs have been widely studied and got tremendous success in different areas. This paper firstly introduces what is GAN, illustrate its structure with 3D visualizations, and discusses the advantages and disadvantages of the original GAN model. Subsequently, this paper introduces some derivative models of GANs, their new features, and their improvements compared with the original model. Finally, this paper summarizes the application fields of GANs, the performance of existing models, some representative works of GAN in image synthesis and editing area, and prospect the possible applications and extensions of GANs in the near future.

A. Introduction to Generative Adversarial Networks(GANs)

Generative adversarial networks (GAN) is a generative model originally proposed by [1]. GAN is one kind of




Fig. 1. The structure of original GAN.

structured learning, it is inspired by the two-person zero-sum game in game theory (i.e. the sum of two-person interests is zero, and the gain of one side is the loss of the other side). Generally, the GAN system consists of a generator and a discriminator(See Figure 1). The generator captures the potential distribution of real data samples and generates new data samples, namely, fake data. The discriminator is a classifier to distinguish whether the input data is real or fake.

As shown in Figure 1, the task for generator G is to train a neural network that is able to convert arbitrary randomly distributed noise, typically called as a latent vector Z , into a synthesized fake data $G(Z)$, and try to make the fake data approach real data x , that is, the training data, as much as possible. At the same time, the discriminator is trained simultaneously as a classifier. Ideally, the data generated by the generator is classified as fake data while the data from the training set as the real data. Therefore, the discriminator is being trained to decide whether the data is real or fake.

In the beginning, the generator generates fake data randomly, therefore, it is easy for the discriminator to identify the fake data. While the generator is optimizing and improving, more and more fake data looks like real data, until the discriminator can't tell which one is real and which one is fake. That means, in the end, the generator can generate realistic data being able to "fool" the discriminator. Therefore, the optimization is done through a min-max loss function, that the "worse" performance of the discriminator results in the "best" performance of the generator. The loss function for discriminator $L_D(D, G)$ is defined as:

$$\min_G \max_D L_D(D, G) = \log D(x) + \log(1 - D(G(z))) \quad (1)$$

where data x is from the real data distribution while latent vector z is from the random distribution (noise). The discrim-

inator D wants to maximize the classification between the distribution from real data $x \sim \text{data}$ and the one from the fake data $x \sim G(z)$. After the optimizations, we hope that $D(x)$ increases which means the discriminator can identify real data more accurately. That is why we maximize D . At the same time, the discriminator hopes $D(G(z))$ decreases, as it wants to identify the fake data more accurately. Therefore, we minimize G . On the other side, as a generator G , we hope that $D(G(z))$ increases which means the fake data $G(z)$ looks like real data after the optimizations. Therefore, the goal of the generator is to maximize the loss function for generator $L_G(D, G)$:

$$\max_G L_G(D, G) = \log(D(G(z))) \quad (2)$$

where similarly, z is a latent vector from a random distribution (noise).

Generator and discriminator can be implemented through deep neural networks which are proved to be robust enough to achieve the goals through a series of experiments. As the trained model is able to generate fake data that conform to the sample distribution of real data without any prior knowledge, therefore, generative models take an important role in unsupervised deep learning to capture the high-order correlation of data without target class label information. By learning the semantic features of the real data, the GAN model can estimate the distribution of training data and generate new data similar to training data.

B. Advantages and Disadvantages of Original GAN

Before the appearance of the original GAN model, there are some existing unsupervised learning-based generative models proved to be efficient for image synthesis. For example, Autoencoder (AE) proposed by [2] is able to convert an input image into a code layer through a neural network (Encoder) and convert such code layer through another neural network (Decoder) back to an image as similar as possible to the original input. During the training process, both encoder and decoder are trained simultaneously and the difference between the input images and the decoded images are backpropagated to the optimizer as the loss function to be minimized. AE and its variational version called Variational Auto-Encoder (VAE) [3] have been successfully applied to image reconstruction and content-based image retrieval [4]. AE overcomes the traditional bottleneck of pixel-based image retrieval approaches by directly comparing the code generated from the encoder instead of comparing the images pixel-by-pixel. Through this approach, AE is able to catch the semantic features of the retrieved images while pixel-by-pixel-based approaches cannot. Another application of AE is image denoising [5]. Before sending the original high-quality input image to the AE, the input image is preprocessed by adding white noises. Then train the AE to generate images as close to the original high-quality input images as possible. Through such a training process, AE is able to reconstruct high-quality images from low-quality noised images input. As another representative generative model, PixelRNN [6] trains the neural network by

estimating the conditional distribution of each individual pixel in a given adjacent pixel (left or upper). In the PixelRNN model, the input of the current pixel and its adjacent previous pixel are both sent into the recurrent neural network as a sequence, the previous pixel is the condition of the current pixel input. Through such a mechanism, the PixelRNN is able to predict another part of an image from one part of the image. Therefore, PixelRNN is proved to be a powerful tool to deal with image completion problems.

Although the existing generative models (such as AE, VAE, etc.) are able to generate new images by itself, the output images are seriously restricted to the training dataset as they are trying to mimic the input image from the training set. Therefore, the synthesized data are very similar to the original input from the existing dataset, and sometimes, such models even directly retrieve images from the dataset. On the other side, PixelRNN is trying to complete the images without considering the global image semantics and trying to combine different images from different catalogs as new images that are semantically meaningless. For example, a car image is completed with a cat face. Therefore, a more intelligent generate model needs to have two advantages: 1) the generated images are semantically meaningful and can be identified as realistic images. 2) the generated images are not retrieved from the dataset and not too similar to the given training samples. With such two advantages, the original GAN model made itself a breakthrough in the state-of-art of machine learning. On one hand, as the GAN has a discriminator to identify the image as real or fake, therefore, if well-trained, GAN is theoretically able to synthesize realistic images, this explains the first advantage. On the other hand, as the generator in the GAN has no input directly from the dataset, therefore, the generator doesn't know how does the realistic data look like and there is no chance for the generator to generate new realistic images by copying the existing images, this explains the second advantage.

However, the original GAN model, inspired by the Nash equilibrium in game theory, still has some inevitable disadvantages. For example, there is no explicit function to evaluate how good the current status of the GAN is. Typically, loss functions can be used to evaluate whether the current neural network is well-trained. If the loss function is close to zero, then the neural network looks "good" enough. But in the original GAN model, the generator hopes to maximize the discriminator's loss function, which means it hopes to fail the discriminator by "fooling" it "smartly". But on the other hand, the discriminator wants to be good at punishing the generator by minimizing its own loss function while maximizing the generator's loss function. Only in that way, the discriminator can be trained to identify the fake data well. Therefore, as a consequence, the optimization will not be guaranteed to converge as their loss functions are adversarial. Also, it is not guaranteed whether the discriminator or the generator is over-trained and it is very hard to be balanced. In the end, there will be either harmonious consistency between discriminator and generator or there is one side over-performed

than another side. For example, when the generator is trying to "fool" the discriminator in some "tricky" ways such as only generate some specific images that can be identified as real images. In that case, the discriminator is not well-trained and the generator will not be able to generate realistic images diversely. For solving such instability, oscillations and divergences issues that potentially exist during the training process of the original GAN model, some variations of the GAN models change the original JS-divergence [7] into other types of divergence such as the f-divergence in the f-GAN [8] or change the distance measure function into Wasserstein distance in the WGAN [9].

II. VARIATIONS OF GENERATIVE ADVERSARIAL NETWORKS

The Generative Adversarial Networks (GANs) was proposed in 2014 and has achieved great influence on the machine learning community. However, the original GAN model is impractical for most of the challenging real-world tasks. Therefore, different variations of GANs models were proposed. For example, Conditional Generative Adversarial Networks (CGAN) [10], Deep Convolutional Generative Adversarial Networks (DCGAN) [11], Information Maximization Generative Adversarial Networks (InfoGAN) [12], Wasserstein Distance-based Generative Adversarial Network (WGAN) [9], etc. Most of these GAN models are milestones that push the original GAN model a great step forward as a powerful tool to solve the real-world problem. In this section, some milestones of the variations of the GAN models will be introduced.

A. DCGAN

DCGAN, short for deep convolutional GAN, firstly proposed by [11], is a milestone in the development of improved GANs model. It combines CNN in supervised learning and the GAN in unsupervised learning, which is robust and convenient for engineering implementation (See Figure 2). In order to enhance the original GAN model, the DCGAN model removes the fully connected hidden layers to construct a deeper neural network. For image analysis, replacing the pooling layer with stride convolutions networks in discriminator and with fractional-stride convolutions in the generator results in a more suitable sampling kernel function for unsupervised learning as such deep neural networks are able to identify the matters in images and capture the features more efficiently. Because of the improved training stability of the DCGAN model, it has been widely accepted and used for academic purposes.




Fig. 2. The structure of DCGAN.

B. CGAN

Inspired by the conditional probability, feeding the GANs with additional conditions results in a novel variation of the GAN, namely, Conditional GAN (CGAN), which is proposed by [10]. In order to solve the problem of how to generate a model given specific requirements, this improved method feeds such conditions (such as labels) into both of the discriminator and the generator networks (See Figure 3). During the training, the generator takes both the label and the latent noise vector z as the input, and pass the label to the discriminator. For discriminator, either the generator generates a poor-quality image or it generates an image not matching with its input label, it will be detected as fake data. In such a way, the discriminator is able to improve both the quality of the generated image and the correlations between the generated image and its label through the feedback for the generator.




Fig. 3. The structure of CGAN.

C. InfoGAN

As proposed by [12], InfoGAN is the GAN extended with information maximizing theory. The output of the original GAN model is $G(z)$, where z is a totally unstructured random vector. During the training, there are no additional restrictions on z , therefore, output $G(z)$ is highly entangled with input z , and make the individual elements in vector z loss their correspondence to the specific semantic features of the data. InfoGAN has made improvements by splitting the original input of the generator model into two parts: (1) the random noise z , and (2) a random latent code c which is used to targeting the semantic features of the data distribution [13]. (See Figure 4).

According to the hypothesis from [12], there should be high mutual information between latent codes c and generator distribution $G(z, c)$, therefore, in order to maximize the mutual information between the c and $G(z, c)$, it incorporates an assistant classifier $Q(c, G(z))$ to classify the latent code c from the synthesized fake data $G(z)$. Therefore, only if c has connections to the salient features of data so that it can be classified correctly by the classifier Q . The discriminator has a boolean output true or false for identifying the real data from fake data, while the output latent code c' are classified through the classifier Q , which is used to be compared with the input latent code c fed into the generator G . If the classified output latent code c is closer to the original input latent code c , it will result in a higher score for the generator G as the generator is able to extract the semantic information from latent code better. As a great milestone, InfoGAN extends the original

GAN model with the ability to synthesize more controllable data representation by tuning the values in c [14], it has taken another step forward in the development of GANs.




Fig. 4. The structure of InfoGAN.

D. EBGAN

According to the original GAN model, in the beginning, the performance of the discriminator is poor and it is improved step by step as more and more real data and fake data are fed in. Therefore, the generator improved extremely slowly in the case that the discriminator is not well trained. However, EBGAN (Energy-based GAN) first proposed by [13] which extends the GANs with an auto-encoder/decoder which is pre-trained (See Figure 5). As the auto-encoder is pre-trained through the real data by minimizing the energy which measures the difference between the decoded images and the input images. Therefore, when the fake images are fed into the autoencoder, they will not be decoded correctly until the generated fake images look real.

By minimizing the energy value, extracted from the well-trained discriminator, which is smaller in the region near the real data domain but higher in the non-real data domain [14], the generator will be improved much faster at the beginning than the original GAN model. Therefore, EBGANs give the original GAN model an explanation of the energy model and generalize the loss function in a novel way so that EBGANs take a further step along with this research direction.




Fig. 5. The structure of EBGANs

III. APPLICATION OF GANs

As a powerful tool to generate realistic data, GANs are widely used in different areas. The most general application of GANs is to synthesize a fake image that looks like a real one and editing the existing image into a synthesized image without distortion. According to the strengths from diverse variations of GANs, GANs has been widely used among different areas to solve real-world tasks and problems. In this section, some of the recent representative works of GANs are selected and introduced.

A. Image Synthesis

Image synthesis is the most general application of the variations of GAN models. There are several famous image datasets serving as benchmarks for evaluating the efficiency of a GAN model. For example, the MNIST database of handwritten digits proposed by [15], CelebA dataset of a large-scale celebrity faces attributes proposed by [16], and Fashion-MNIST proposed by [17]. Different variational GAN models demonstrate




Fig. 6. Performance of DCGAN: (a) samples from the MNIST dataset, generations from the original GAN, and generations from the DCGAN. Image is cited from [11].

the improvements in their performances on the benchmark database. [11] devised DCGAN to generate different ranges of images including bedroom and face images, and prove that the synthesized images are not retrieved from the existing image database. Its performance compared with the ground truth of MNIST and the original GAN are shown in Figure 6. Beyond this baseline, CGAN is moving a step forward which is able to generate the samples with specific conditions as shown in Figure 7. As shown in Figure 8, InfoGAN is able to extract the semantic information from the sample images and find the relations between the latent code c and such semantics. This provides with users a more controllable interface to generate the desired images. Based on the performance of DCGAN, EBGAN makes further improvements in generating high-quality images. As shown in the LSUN bedroom dataset (See Figure 9), the left images are generated with the DCGAN, the right images are generated with the DCGAN, some noise can still be found in the DCGAN, but in EBGAN, the noises are less obvious.




Fig. 7. Generated MNIST digits with CGAN, rows (from top to bottom) represents conditioned labels from "2", "3", "4", "5", to "6". Image is cited from [10].




Fig. 8. InfoGAN's latent code c has image semantics: In (a), varying c_1 corresponds to different digit types; in (b), a small value of c_2 denotes left-leaning digit whereas a high value corresponds to right-leaning digit; in (c), c_3 smoothly controls the width of the strokes. Image is cited from [12].

B. Image Editing

As GAN models and technologies are getting mature, the extensions of the GANs are increasing rapidly, especially in the area of image editing. As traditional ways to edit

the image, the popular image processing software, such as Photoshop developed by Adobe company, is widely used among multimedia artists and designers. However, even the software is getting more powerful and provides more easy-to-use functions, there still need lots of manual efforts from the designers. Therefore, GANs' impressive synthesis ability attracts more and more researchers to develop image auto-editing tools through deep neural networks. Typically, image editing is divided into two types: local editing and global editing. The most famous application of deep learning on global editing is the image style transfer technique [18]. Most of the existing global editing algorithms are built on image-to-image translation networks and result in surprising results [19]. However, unsupervised image-to-image translation and image local editing remain challenging before the appearance of the GANs models. As local image editing techniques require more conditions than the global image editing (such as whether the stitching is natural, whether the edited part is consistent with the full image content, etc.), this challenges of the local image editing techniques provide a perfect stage for the GANs to demonstrate their specialty. As the GANs




Fig. 9. This figure compares the generator performance between DCGAN (a) and EBGAN (b). Image is cited from [13].

model can repair and complete the missing area of the image according to its surrounding area from the semantic level, it provides more natural and acceptable results than the nearest neighbor stitching method. Typically, GANs work through the Context Encoder (CE) [20], which are widely used for high-resolution image patching under specific circumstances. CE includes two parts: Encoder and Decoder, which typically incorporate DCGAN applied with ADAM optimizer to achieve image auto-completing with highly satisfying quality. For example, [21] successfully change the status of the people's eyes in an elegant manner using the Exemplar GAN, a variation of CGAN. As shown in Figure 10, Exemplar GAN can achieve photo-realistic, high-quality, and personalized inpainting function to edit the image in an astonishing way. As another impressive work for image completion with high consistency between local image and global image, high image resolution, and extremely photo-realistic image content, [22] successfully achieved the real-time and interactive auto-completing of local image editing skills using GAN model (See Figure 11). At the same time, the face editing skills also improve. As shown in the Figure, [19] it successfully devised a conditional DCGAN-based approach to automatically edit the human face and hair image in a time-saving manner simply by sketching on the face image with simple curves. This is a great invention as it opens the eyes for people towards photo editing through sketching, the most convenient, and straightforward way to edit the images. Also, there are many other types of

applications of the GANs that are not listed here, which will be briefly discussed in the discussion sections.



Fig. 10. Eye editing using Exemplar GAN. Image is cited from [21].



Fig. 11. Auto-image completion using GAN. Image is cited from [22].



Fig. 12. Sketch-based interface for face image auto-editing using CDCGAN. Image is cited from [19].

IV. DISCUSSIONS AND CONCLUSION

Besides image synthesis and editing, GANs has more various kinds of applications among different research areas, for example, GANs have deep influences on image enhancement, super-resolution. The same kind of image, such as face image enhancement from low-resolution face image to high-resolution image, can be better implemented by GANs. The idea is to use a low-resolution image as a constraint condition to generate the realistic high-resolution images [11]. As another example, the super-resolution energy adversarial network (SRGAN) [23], without depending on refined data sets, it can enhance and denoise various types of images. The features of SRGAN includes: (1) completing the missing part of images with the constraint on the global context of the generated image to ensure the smoothness of the result; (2) SRGAN feed the generated data and real data into the VGG-19 network respectively, define the loss items according to the difference of the feature map and adds normalization to the output [14]. Finally, after combining these loss items including mismatching loss, image smoothing loss, and feature map difference loss, the super-resolution image is generated by feeding into the GANs framework through minimizing the loss functions.

Besides, the GANs have been extended into the sketch restoration and colorization techniques. Sketch restoration refers to the technique that can convert the drawing of sketches into realistic color images. Its special case is portrait restoration. A DCGAN can complete the restoration of draft drawings [24]. it gives the training methods of two kinds of image databases with corresponding relations, finds the relationship between database images, and gives a general solution of sketch-to-image conversions, such as the conversion of the street scenes, the sketch of buildings, the terrain elevation from a contour map. Traditional image coloring methods do not use massive data, and image coloring can not be applied to all

types of pictures. [13] use Patch GAN network to complete the picture coloring work, and make full use of massive data to achieve better results. This proves that GAN plays an active role in sketching techniques.

In conclusion, generative adversarial network(GANs) has become one of the most important and influencing methods in deep learning. It has the advantages of fully fitting data, faster synthesis speed, and realistic data generation. The academic research of the GANs model is progressing rapidly. The original GANs model is trained by MinMax optimization. The conditional generative adversarial network(CGAN) adds preconditions to input data in order to provide the controllability of the output. DCGAN, a deep convolution generation adversarial network, proposes a stable training network structure to prevent training collapse. InfoGAN controls semantic change through latent code and extracts the relations between the latent code and its corresponding semantic feature of the training data. EBGAN explains the adversarial network from the perspective of the energy model. WGAN defines a more smooth distance measure for the loss function and gives a better mathematical definition of the distance between real data distribution and fake data distribution, which theoretically solves the problem of training collapse and instability. As a powerful, robust, and reliable tool, the GANs model has been widely used in image synthesis, image editing image repair, image denoising, sketch restoration, sketch colorization, and other image processing area. At present, however, there is no quantitative standard for evaluating the realism of synthetic images, that is, it is hard to quantitatively measure how realistic an image is. It can be only judged subjectively whether the synthesized images look natural and realistic. This remains an open topic for the GANs community. In the near future, GANs will be applied to more general applications not only within the image processing domain. For example, music synthesis using GANs to generate different styles of music composed by different composers or even create innovative music styles in a "masterpiece" level like famous musicians such as Beethoven and Mozart. The GANs can also be extended to write articles and poems with different styles and cultural backgrounds. When the GANs are getting more mature, even the GANs can take over human jobs such as movie directors who are directing new movies including story design and character artistic design, choreography artists who are designing new dance for a group of dancers and so forth. In one last word, GAN is going to change human lives from different aspects in the near future.

V. ACKNOWLEDGMENT

I acknowledge my directed research course advisor Dr. Jana Kosecka, this review paper has been completed under her instructions and supervision. This review paper has been inspired by one of her publication: *Introduction to Multiview Rank Conditions and their Applications: A Review*.

REFERENCES

- [1] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [2] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *science*, vol. 313, no. 5786, pp. 504–507, 2006.
- [3] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [4] A. Krizhevsky and G. E. Hinton, "Using very deep autoencoders for content-based image retrieval," in *ESANN*, vol. 1. Citeseer, 2011, p. 2.
- [5] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol, "Extracting and composing robust features with denoising autoencoders," in *Proceedings of the 25th international conference on Machine learning*, 2008, pp. 1096–1103.
- [6] A. v. d. Oord, N. Kalchbrenner, and K. Kavukcuoglu, "Pixel recurrent neural networks," *arXiv preprint arXiv:1601.06759*, 2016.
- [7] B. Fuglede and F. Topsøe, "Jensen-shannon divergence and hilbert space embedding," in *International Symposium on Information Theory*, 2004. ISIT 2004. Proceedings. IEEE, 2004, p. 31.
- [8] S. Nowozin, B. Cseke, and R. Tomioka, "f-gan: Training generative neural samplers using variational divergence minimization," in *Advances in neural information processing systems*, 2016, pp. 271–279.
- [9] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein gan," *arXiv preprint arXiv:1701.07875*, 2017.
- [10] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014.
- [11] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.
- [12] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel, "Infogan: Interpretable representation learning by information maximizing generative adversarial nets," in *Advances in neural information processing systems*, 2016, pp. 2172–2180.
- [13] J. Zhao, M. Mathieu, and Y. LeCun, "Energy-based generative adversarial network," *arXiv preprint arXiv:1609.03126*, 2016.
- [14] A. Odena, C. Olah, and J. Shlens, "Conditional image synthesis with auxiliary classifier gans," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70. JMLR. org*, 2017, pp. 2642–2651.
- [15] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [16] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proceedings of International Conference on Computer Vision (ICCV), December 2015*.
- [17] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *arXiv preprint arXiv:1708.07747*, 2017.
- [18] L. A. Gatys, A. S. Ecker, and M. Bethge, "A neural algorithm of artistic style," *arXiv preprint arXiv:1508.06576*, 2015.
- [19] T. Portenier, Q. Hu, A. Szabo, S. A. Bigdeli, P. Favaro, and M. Zwicker, "Faceshop: Deep sketch-based face image editing," *arXiv preprint arXiv:1804.08972*, 2018.
- [20] S. K. Yelamarthi, S. K. Reddy, A. Mishra, and A. Mittal, "A zero-shot framework for sketch based image retrieval," in *European Conference on Computer Vision. Springer*, 2018, pp. 316–333.
- [21] B. Dolhansky and C. Canton Ferrer, "Eye in-painting with exemplar generative adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7902–7911.
- [22] S. Iizuka, E. Simo-Serra, and H. Ishikawa, "Globally and locally consistent image completion," *ACM Transactions on Graphics (ToG)*, vol. 36, no. 4, pp. 1–14, 2017.
- [23] S. Lian, H. Zhou, and Y. Sun, "Fg-srgan: A feature-guided super-resolution generative adversarial network for unpaired image super-resolution," in *International Symposium on Neural Networks. Springer*, 2019, pp. 151–161.
- [24] A. B. L. Larsen, S. K. Sønderby, H. Larochelle, and O. Winther, "Autoencoding beyond pixels using a learned similarity metric," *arXiv preprint arXiv:1512.09300*, 2015.