



A COMPREHENSIVE STUDY ON VARIOUS ASPECTS OF GENERATIVE ADVERSARIAL NETWORKS (GANS)

Avi Maheshwari¹, Akash Kumawat², Aakash Gupta³, Dr. Sunil Srivastva⁴

^{1,2,3}Students, JECRC Jaipur, India

⁴Associate Professor, JECRC Jaipur, India

ABSTRACT:

GANs have emerged as one of the foundational frameworks of contemporary generative learning, providing state of the art solutions for generation as well as data augmentation. In this detailed work, we explore the specifics of GANs to include a description of their raw ideas, structural complexities, and the training of the adversaries that allow the GANs to work. At the center of GANs is the learning process of the generator and discriminator, a pair that is engaged in a principled form of deceit concluding in the formula for the loss function that defines the learning process. The paper also discusses other types of GANs which are a small subset of GAN classes that have been uniquely designed to address various hurdles or improve the performance of these models in particular tasks. Starting from image generation to more bandwidth and advanced domain adaptation tasks, GANs have established their supremacy. But we also know that the road to full attainment of the capabilities of GANs is not without some challenges; this paper also discusses the challenges and limitations within GANs such as training instability and mode collapse problems, and possible resolution to them. At the end of the paper, the authors discuss possible future trends in GAN research and development, underlining the requirement of better stability, targeted, and explainable models. In an attempt to present the concept of GANs in simple terms, this research paper will seek to avoid the use of overly technical language to explain the concept of the technology to readers, as well as to general readers interested in knowing more about this type of self-supervising unsupervised learning algorithm so as to be useful reference body of information for anybody who wishes to get acquainted with or work on this technology in any possible following advancements.

Keywords: Generative Adversarial Network, Unsupervised Learning, Generator, Discriminator, Convolutional Neural Network, loss function, biases, weights.

[1] INTRODUCTION

Gen is short for Generative, while the second term, Adversarial Networks, or GANs, refers to a specific type of Artificial Intelligence that is nested within the even larger

category of machine learning. They were developed within the unsupervised learning regime which locally allows the system learn about the patterns and representations of the data without specific instructions on how to do that. The idea of GANs centered on a generator and a discriminator or adversarial networks working in tandem and being trained at the same time. During training, a competitive min-max game unfolds: While the generator wants to maximize the discriminant loss it tries to maximize the probability of making more errors, on the other hand, the discriminant loss tries to minimize this probability.

Text Box The generator’s function is to generate data that are not differentiable from real data. From here it obtains the capacity of transmutation from a learned space to the data distribution desired. The discriminator, for the other part, processes the information that it gets; the real data and the fake data created by the generator. Based on GAN, the generator’s measure of deception is to generate fake samples while the discriminator’s strategy is to distinguish them. Still, in Generative Adversarial Networks, it is possible to notice a kind of teamwork because two parts of a computer program are involved. This one is like an artist and is making new pictures or data and is called a generator. The other part behaves like a validator or the discriminator which aims at determining whether the pictures are invented or real. He has to keep on drawing better and better picture and it helps when the judge has to make decisions. This continues until the picture making apparatus is created to be so precise that the pictures produced are indistinguishable from the real ones.

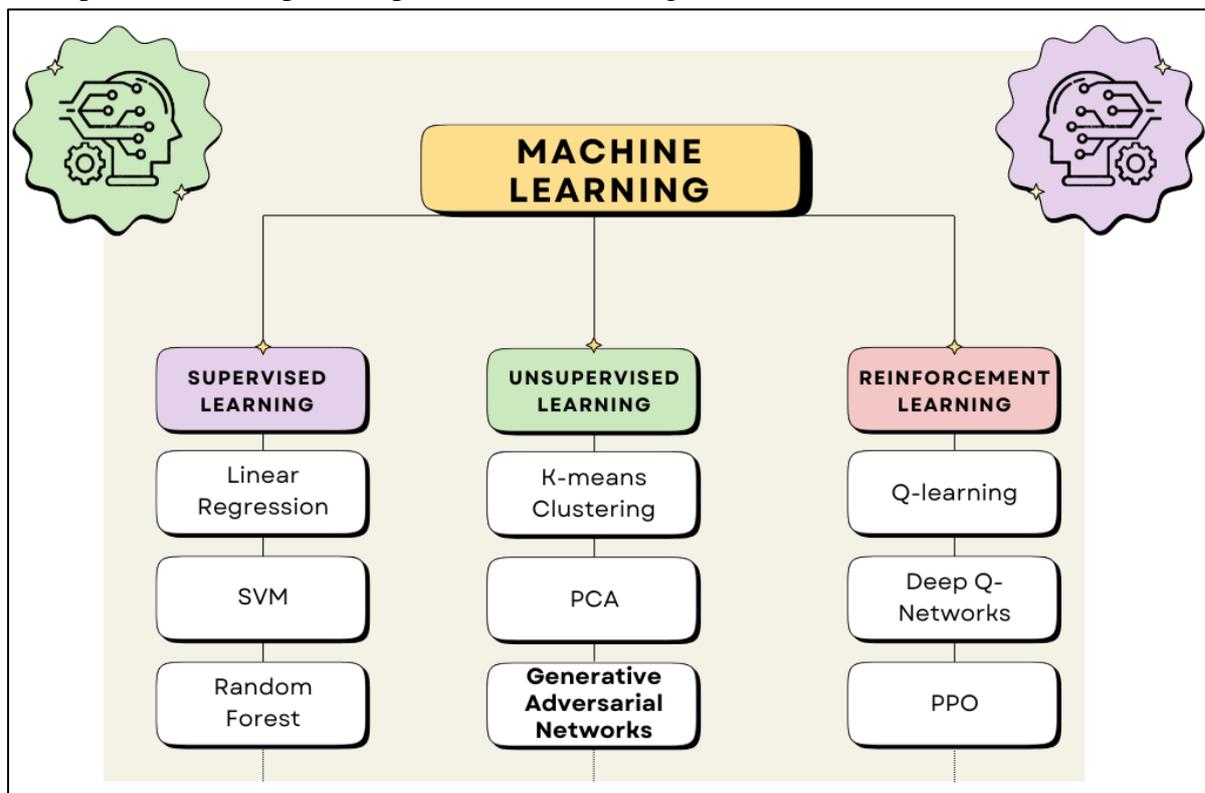


Figure 1 Classification of Machine Learning

2. BACKGROUND

2.1 Basic Architecture

The fundamental architecture of GANs consists of two neural networks that engage in a game-theoretic confrontation: the generator and the discriminator.

Generator: In the world of GANs, the generator is a neural network that is an artist, indeed generating novel data samples similar in distribution to the data being generated. The generator's primary purpose is to learn how data is distributed in the training data set and, based on that knowledge, synthesize new instances that can belong in the same distribution. This is achieved by transforming a point from a latent space step, usually a random vector, to the data space. For instance, in image generation scenarios, the generator models convert random noise vector into an organized image that looks similar to the images on the training set.

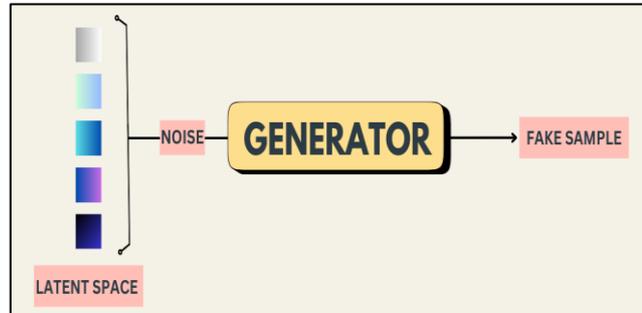


Figure 2 Working of Generator

To do this, the generator leverages several layers that enhance the up sampling of the input latent vector into the equal output dimension. The generator initially produces a layer, which reapplies and remodels the input latent vector for the following transposed convolutional layers. These then upscale the feature maps and every layer add more details into the former. The commonly applied operation after each transposed convolution is batch normalization, which maintains stability in the learning process and enables the network to learn the mean and variance of each layers' output. Whereas non-linearity is an essential capability of the network that is required when modeling the data distribution, activation functions add such non-linearity into the neuron.

Here's a detailed breakdown of the typical layers and their roles in the generator's neural network architecture:

1. Input Layer: The first layer of the generator takes an initially random vector hence the term latent space or z-vector usually from a Gaussian distribution. As an illustration, this vector defines the first data generated in an advertising application.

2. Fully Connected (Deep Connected) Layers: It is worth mentioning that, after the input layers, the generator includes the fully connected layers. These layers help to reduce the dimensionality of the latent space and project it into a shape that fits further convolutional layers. The neurons within these layers are Highly interconnected.

3. Convolutional Layers: These layers are the ones most important for the generation of image data. It involves learnable filters that enable the detection of spatial features in the input data. Along with convolutional layers, other methods such as upsampling are applied in order to increase the feature map size.

4. Transposed Convolutional Layers: These are also referred to as deconvolutional layers

They act in the reverse way as convolutional layers. They resize the feature maps, widen, and heightening them while possibly thinning in depth. This is done through the principle of scaling a single input value across multiple output planes.

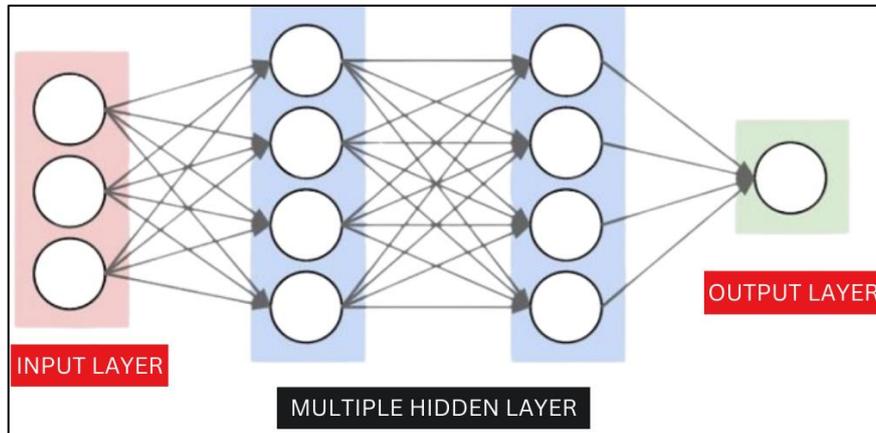


Figure 3 Artificial Neural Network

5. Batch Normalization: This has the effect of steadying and accelerating training in deep neural networks. It scales the output of a preceding activation layer to decrease the mean and increase the standard deviation according to the batch data. It also aids in controlling omissions resulting from internal covariate shift during training.

6. Activation Functions: Non-Linear Activation functions are used in Neural Networks because they are capable of capturing time-related changes. In the generator, the activation function used in all layers are ReLU (Rectified Linear Unit) or LeakyReLU. The output layer often utilizes the tanh activation function, as it scales data in the range of $[-1, 1]$, favorable for image data.

7. Output Layer: It is the last layer to be described in generators, and its output is the synthetic data instance. In image generation tasks, this layer outputs an image of the required height, width, and the number of channels in the image for colored images it could be 3.

Discriminator: The discriminator inside of a GAN is a neural network that serves as the key way by which the GAN learns. It is mainly used for the classification task which means to decide whether the input data is real or fake data generated by the generator. Since the discriminator plays a significant role in determining the quality of the data produced by the GAN, the mechanism is a central aspect of the adversarial learning process. First, the discriminator assesses the credibility of the data it receives which may be from the training set that is real data or from the generator which is fake data. It provides an estimate for each data instance regarding how likely they are to be real. The

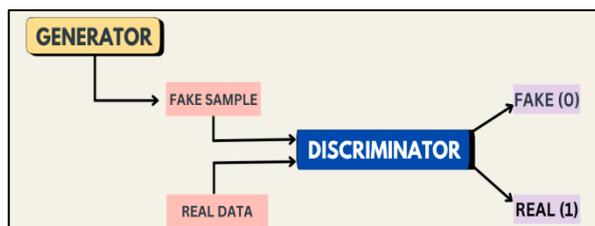


Figure 4 Working of a Discriminator

discriminator tries to distinguish between true and false data to achieve the highest accuracy. Its architecture elements serve to denote whether the inputs used in the training set are real or fake.

Here is a detailed description of the typical layers found in a discriminator's neural network:

1. Input Layer: The input layer works as the reception layer for the discriminator network of the model. Namely, the mentioned layer—Image Discriminator—takes images from the data set or fake images created by the generator. Size of the image taken as the input layer is equivalent to the size of the images processed.

2. Convolutional Layers: When working with image/data, convolutional layers are in the forefront of discriminator's construction. These layers linearly transform the input and produces feature maps to capture spatial hierarchies. Moreover, convolutional layers can also use stride to get rid of the spatial dimension of the feature maps, in this way, downsampling occurs.

3. Activation Functions: In particular, in conventional CNNs, an activation function is often used to process the data after each convolutional layer. When applied to the discriminator, the non-saturated activation function used is the LeakyReLU function. LeakyReLU enables a small nonzero gradient throughout the training process in case the unit is non-active.

4. Batch Normalization: Batch normalization can be used after convolutional layers to make the learning process less volatile. It operates on activation and scales it to bring the output of a layer closer to standard mean. It can be faster the training process and enhance the performance of the network in general.

5. Convolutional Layers: The discriminator, in turn, starts with several convolutional layers to extract feature maps. These layers transform the input data (in general images) by using filters to recognize patterns including edges, textures and shapes. The next layers represent even higher layers of features that are captured by the model.

6. Strided Convolutions: Unlike in the past, contemporary GANs resort to strided convolutions while pooling layers are not employed. Stride reduces the spatial dimensions of feature maps as it skips over a number of pixels in the image during convolution. This allows getting down sampling, which prevents the model from having dedicated pooling layers.

7. Fully Connected Layers: The last section of the discriminator structure involves fully connected layers wherein the convolutional layer features are merged. These layers reduce the spatial dimensions into one the model receives and makes a final decision from the combined inputs.

8. Output Layer: The output layer, consisting of a single neuron, activates by the sigmoid function. This neuron coded the input data as the probability of the input being real rather than a fake image created by the generator. Probability A reveals the realistic input with high probability, while probability B is low showing the generated sample.

2.2 Adversarial Training Process

It starts with the initializations of both the networks with random weights and biases being assigned to the text feedbacks. The generator generates data from the noise vector and the discriminator compares it with the real data samples. The discriminator employs a loss function to bring about an improvement in the ability of the model to separate real data from fake data. On the other hand, the generator tends to minimize this loss which in return improves the quality of the artificial data that it produces.

In the discriminator's training phase, it receives both authentic data and artificial data created by the generator. The discriminator evaluates the genuineness of each data item and assigns a likelihood of it being real. The discriminator's loss is computed using the binary cross-entropy loss function, which quantifies the difference between the predicted and actual classifications. The weights of the discriminator are updated via backpropagation, a technique that computes the gradient of the loss function with respect to the network's weights, facilitating the modification of weights to minimize the loss. It's important to note that the generator's weights remain unchanged during the training of the discriminator.

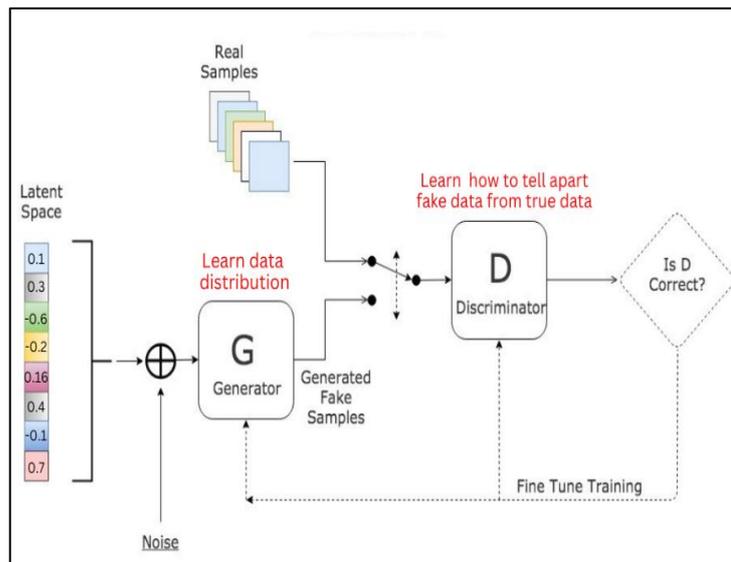


Figure 5 Adversarial Training Process

Similarly, the generator is trained to produce input data that discriminator should perceive as genuine data. The generator loss is also calculated using the binary cross-entropy function, which points to how well it does in deceiving the discriminator and causing it to fail. Most of the weights of the generator are updated through the backpropagation process feedback from the discriminator. Here it should be mentioned that through this stage, the weights of discriminator remain fixed.

This training process is slowly iterated for numerous epochs until both networks are optimal in their output. The generator over time gets better at generating data that is more realistic while the discriminator receives training on how to diagnose data sets. In GAN training process, there are two stages, in which one-stage will incorporate the training of discriminator and then the generator. Firstly, the discriminator learns the real and false samples in order to maximize discriminative ability of the classifier. After that, the generator is trained to minimize the same loss function so as to improve its capability to generate data that the discriminator will classify as real (Real). This is done in multiple epochs where each network improves the performance of the other in each cycle.

The state of the generator model is saved once performance has achieved an adequate level. This make the generated data to be repeatable or even fine tune at a later stage if needed. They often involve the preservation of the weights and biases obtained during the course of the Generator's neural network, if the learned distribution must be used in future.

[1] Backpropagation is a generalized strategy familiar to everyone who deals with neural networks, including GANs. It is the various derivatives of the loss functions that quantify

how the model parameters should be adjusted in a bid to improve the fitness of the model. The only drawback is that during minimizing the losses of the discriminator and generator networks backpropagation is involved. The parameters of the discriminator only change when the discriminator is being trained and the same applies for the generator. This ensures that each of the networks is trained committed to a fixed rival.

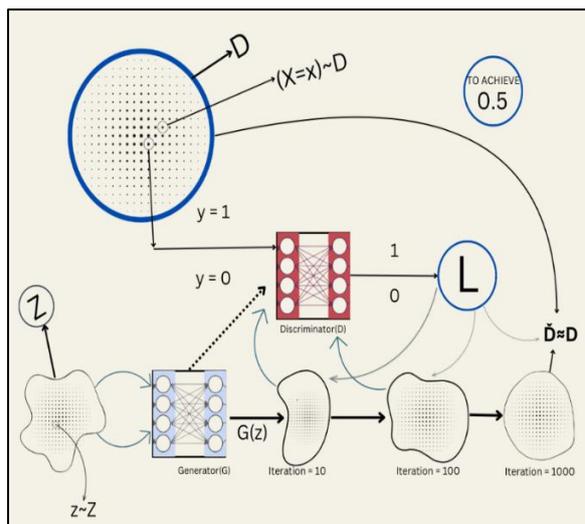
In the case of GANs, convergence is achieved when the discriminator is no longer able to distinguish between real data and synthetic data generated by a certain percentage greater than chance. This implies that the generator is able to generate data points that is as realistic as ‘real live human being’. However, it is not easy to reach that level of convergence since the training process is dynamic and there might be such problems like mode collapse. However, another point at which convergence is achieved, they suggest that the adversarial-training has been beneficial and the generator should be capable to generate realistic fake data.

3. MATHEMATICAL FORMULATION

In Generative Adversarial Networks (GANs), there exists a training process referred to as the Adversarial Training Process that makes it possible to synthesize points of data that cannot be clearly distinguishable from real data.

First, we sample a noise variable z from prior distribution of noise which is represented as, Z It is a noise input that goes into the generator network represented as $G(z)$.

The generated data is fed to discriminator network, and also real data sample is passed to the same discriminator network. The discriminator {a binary classifier} assigns labels to the inputs: Therefore, $y=1$ for real texts testing (Real Sample) and $y=0$ for testing with synthetic text data being tested (Fake Sample). The model is supposed to receive and provide its inputs and return Real or Fake in this case.



Here in this Figure 6:
D = MNIST data set
X = MNIST data set examples
x = instantiation of X
G = generative network
Z = random distribution (random probability distribution)
z = sample of Z
G(z) = generated distribution by the generator
 \hat{D} = Nearest Generated Distribution.

Figure 6 Representation of adversarial training process

When the discriminator successfully classifies the instance, the loss function is calculated and the error was propagated through the Generator networks. This error means that the

generator was unable to deceive the discriminator that it produced realistic images, and so, the weights of the generator are updated.

When the generator's weights and bias are updated, it samples again from the noise distribution, and as the iterations go on the $G(z)$ will look more like the real data distribution. This process enhances the created distribution gradually until the discriminator finds it difficult to differentiate a fake and original data.

If the generated distribution becomes similar to the real (original) data distribution, then the discriminator is unable to correctly classify and thus the decision value hovers around 0.5 making the discriminator optimal.

At this point, the generator is merely deceiving the discriminator, making it impossible for it to differentiate between fake and real. The final generated distribution is \check{D} and it is similar to the actual data set distribution D and the data points of the \check{D} are synthesized data points and represented as $G(z)$. This is the final process in the training process of GANs that set them apart as generative networks.

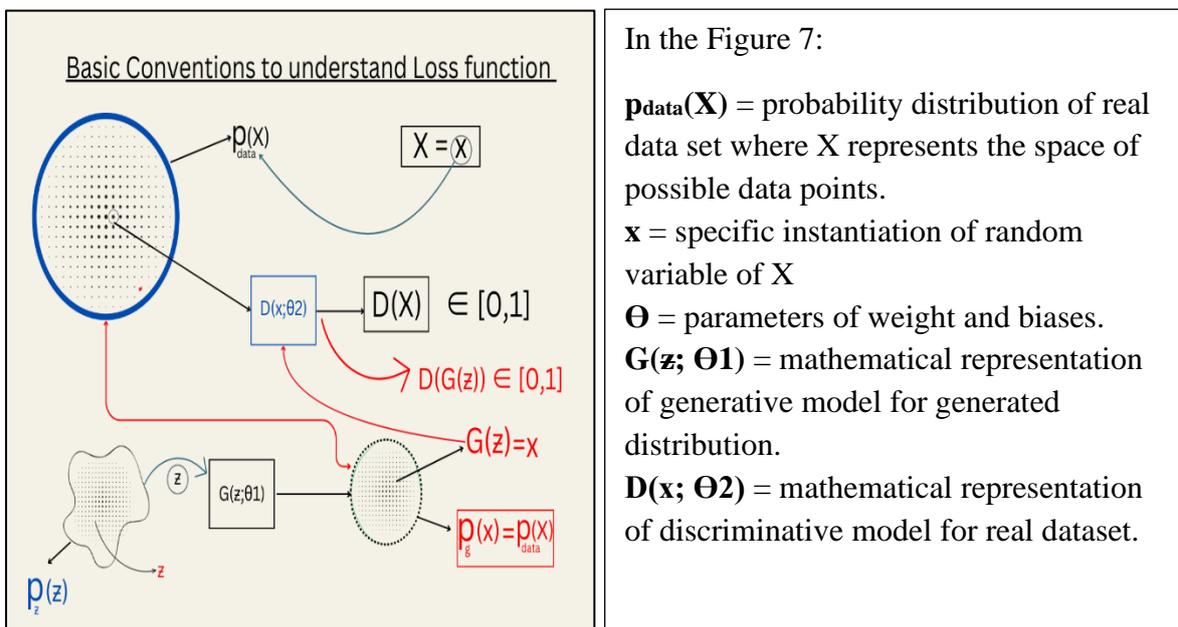


Figure 7 Basic Conventions to understand Loss function

This input the $D(x; \Theta 1)$ results the transform variable which is $D(X)$ this symbol shows what probability of data X received from real ground, normally its range of value is between 0 to 1 only.

Generator's distribution →

We also have the probability distribution function $p_z(Z)$ for samples Z which are enlisted from the noise terms. after we feed the z to the generator which is parameterized by $G(z; \Theta_2)$. G and D are function values in which the above differentiation process will be done. Finally, only the generator can approximate such a distribution as would be in the actual data set at that number of iterations equal to the data set. And this fake distribution is denoted as $p_g(x)$.

And the more we move forward through the generator, if we introduce, for example, a probability variable z , this generator will transform it into a new form $G(z)$, and this form of space is called $p_g(x)$.

Our aim to make $p_g(x)$ as close as to $p_{data}(x)$

So

$$P_g(x) \approx P_{data}(x)$$

And $G(z) = x$

Now when we fed the $G(z)$ to the discriminator it transforms into $D(G(z))$ which represent the probability.

By loss function Equation (Binary cross-entropy) →

$$L(\hat{y}, y) = y \cdot \log \hat{y} + (1 - y) \cdot \log (1 - \hat{y})$$

\hat{y} = output from above neural network, means the reconstructed data or the reconstructed value which we are getting at the end.

y = original data in which the analysis is to be made at the input of the neural network

The data originating from the probability distribution $P_{data}(X)$ is labelled as $y=1$ and the predicted label (\hat{y}) is given by the discriminator function $\hat{y}=D(X)$

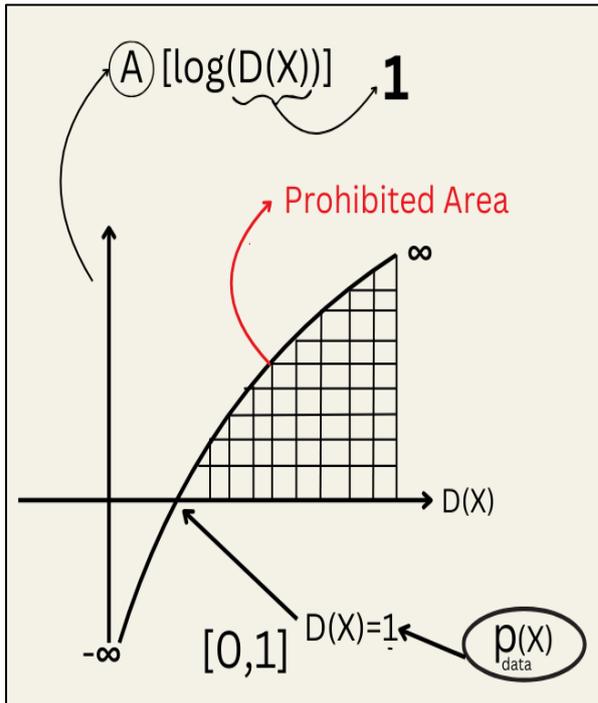
So putting this we obtain

$$L(D(X), 1) = \log(D(X)) \quad \text{-----(A)}$$

Now the data coming from generator the label is $y=0$ and $\hat{y}=D(G(Z))$

Putting the above values in Loss Functions

$$L(D(X), 1) = \log (1 - D(G (Z))) \quad \text{-----(B)}$$



Graph 1 Maximization of Equation A for Discriminator

Indeed, the discriminator's goal is to accurately distinguish between fake vs real data set, so for the equation A and B should be maximized.

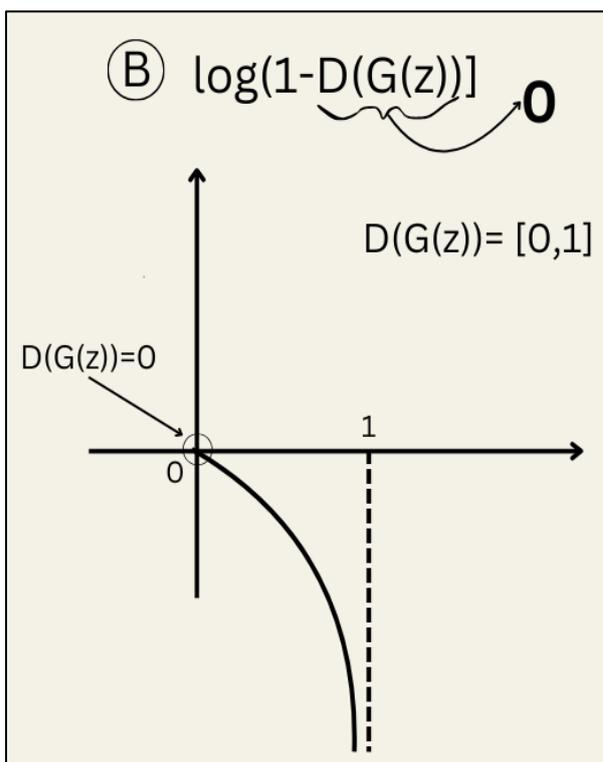
Maximization of eq. A:

$$L(D(X),1) = \log(D(X))$$

$D(X)$ =The probability that X originates from real database.

We get max of the curve:

$\text{Log}(D(X))$ at $D(X)=1$



Graph 2 Maximization of equation B for Discriminator

Maximization of eq. B:

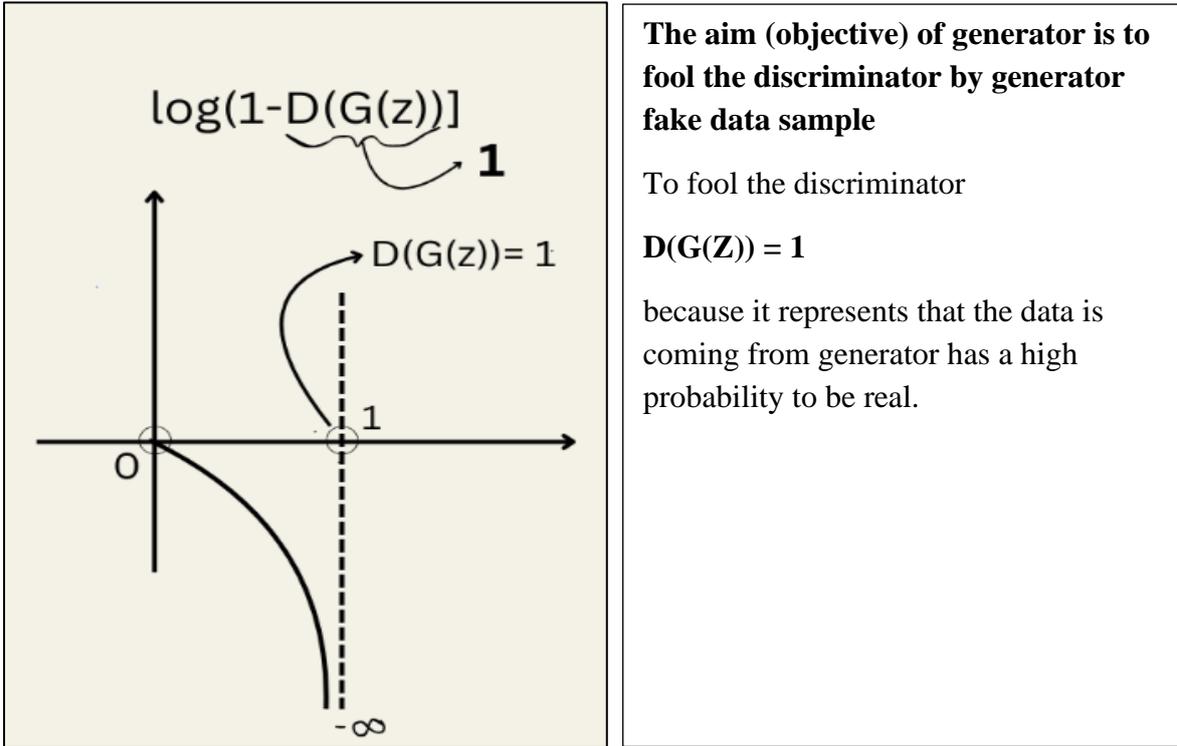
We get max of curve at

$\log(1-D(G(Z)))$ at $D(G(Z)) = 0$

Now maximization eq. of discriminator is

$\text{Max} \{ \log(D(X)) + \log(1-D(G(Z))) \} = D$ (discriminator model)

For generator →



Graph 3 Maximizing Probability of Generator's Output

Now the eq. for discriminator and generator is ->

$$\min_G \max_D \{ \log(D(X)) + \log(1-D(G(Z))) \}$$

To consider all the sample from real dataset we take expectation of above eq.

$$\min_G \max_D V(G, D) = \{ E_{x \sim p_{data}(x)} [\log D(x)] + E_{z \sim p(z)} [\log(1-D(G(z)))] \}$$

Equation Given by **Ian J. Goodfellow**

4. VARIANTS OF GANS:

1. Vanilla GAN:

The original GAN, or Vanilla GAN, involves a generator (G) and a discriminator (D) playing a minimax game with the value function:

$$V(D, G) = E_{x \sim p_{data}(x)} [\log D(x)] + E_{z \sim p(z)} [\log(1-D(G(z)))]$$

This basic form has been the foundation for many subsequent variations.

2. Conditional GAN (CGAN):

Both in Vanilla GAN and CGAN, there is a generator and a discriminator but in CGAN, the generator and discriminator are conditioned on some extra information, say y , such as class label. The objective function becomes:

$$V(D, G) = E_{x \sim p_{data}(x)}[\log D(x|y)] + E_{z \sim p_z(z)}[\log(1 - D(G(z|y)))]$$

This allows the model to generate data specific to the given condition.

3. Deep Convolutional GAN (DCGAN):

Much improved over the original GAN, DCGAN incorporates convolutional layers in to the generator and discriminator, thus providing higher image quality. The architectonic rules are similar to the batch normalisation and strided convolution, while the chief loss function resembles the original GAN.

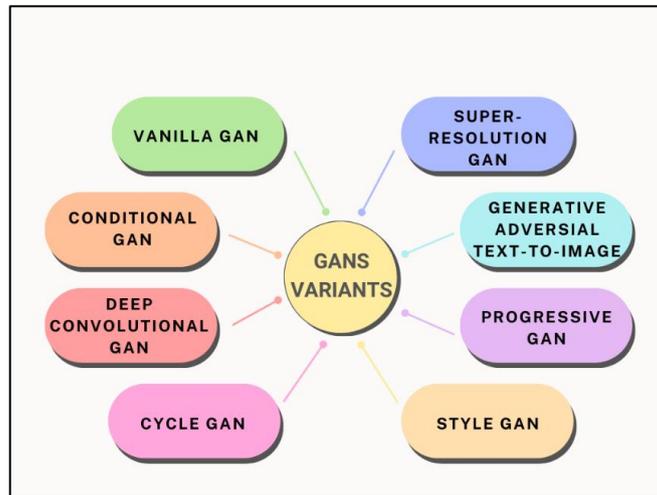


Figure 8 Variants of GANs

4. CycleGAN:

CycleGAN is designed for unpaired image-to-image translation, using two generators and two discriminators. It introduces a cycle consistency loss to the GAN objective, ensuring that an image can be translated from one domain to another and back again without loss of content:

$$V(D, G) = V_{GAN}(D, G) + \lambda \cdot V_{content}(G)$$

$$V(D, G) = V_{GAN}(D, G) + \lambda \cdot V_{cycle}(G, F)$$

where V_{cycle} is the cycle consistency loss and λ is a weighting parameter.

5. Progressive GAN:

Another extension of standard GANs is based on the ‘progressive growing’ concept, where the model learns at first low resolution images and gradually increases the resolution

through further layers of the networks. This does not modify the structure of the base GAN equation but modifies the way in which the training process is conducted with the aim of enhancing stability and images.

6. Super-Resolution GAN (SRGAN):

SRGANs are used for super-resolution tasks. They introduce a perceptual loss function that combines content loss and adversarial loss to produce high-resolution images from low-resolution inputs:

Where V_{content} is the content loss, typically a mean squared error between feature representations of high-resolution and generated images.

7. Text-to-Image GAN:

The Text-to-Image GAN models created images from textual captions. They often employ a text encoder to transform words into the model's input space and rely on this embedding when generating images. The objective function depends on the textual representations, which are analog to the class labels in CGAN.

8. StyleGAN:

StyleGANs generate images with controllable and fine-grained styles. They introduced the concept of style mixing, where different layers of the generator are influenced by different parts of the input noise vector. The objective function includes a style-based generator architecture that allows for precise control over the generated image's style.

5. APPLICATION

1. Generating Realistic Data for Training:

- **Limited Data Scenarios:** When real-world data collection is expensive or scarce, GANs can generate synthetic data that closely resembles the actual data distribution. This is particularly useful in medical imaging or scientific simulations where obtaining real data might be limited.

2. Image and Video Synthesis:

- **Realistic Image Generation:** This can be used for art creation and animations, advertisement images, video game graphics, etc.
- **Image Editing and Enhancement:** Tasks like photo restoration, removing noise from images, or increasing image resolution can all be tackled by GANs.
- **Video Prediction and Generation:** GANs can be used to predict future frames in a video sequence or even generate entirely new videos that follow a specific style.

3. Image-to-Image Translation:

- **Style Transfer:** GANs can be used to translate images from one style to another, such as converting black and white photos to color or translating sketches into realistic images.
- **Medical Imaging:** Segmentation of medical images for disease diagnosis can be improved by using GANs to generate synthetic data with specific pathologies.

- **Satellite Image Analysis:** GANs can translate satellite images into more readily interpretable formats like Google Maps, aiding tasks like urban planning and disaster management.

4. Other Applications:

- **3D Object Generation:** Creating realistic 3D models from 2D images or other data is a developing application of GANs. This can be handy when relying on virtual reality, computer games, and computer-aided design.
- **Data Augmentation:** For tasks like text-to-speech or speech synthesis, GANs can be used to generate additional training data, improving the performance of machine learning models.
- **Security Applications:** Understanding how GANs can be used to generate adversarial examples (crafted inputs that fool machine learning models) helps improve the robustness of these models in security-critical applications.
- **Anomaly Detection:** GANs can learn the distribution of normal data and, as a result, can be used to detect anomalies or outliers in datasets, which is valuable in fraud detection and monitoring systems.

6. CHALLENGES AND LIMITATIONS

1. Mode Collapse: Some of the following issues remain problematic in the current use of GANs within the field; Mode discard; this is a potential problem whereby the generator is able to minimize data variation to a limited range of outputs. This occurs when the generator starts at the output or a limited set of outputs at the start of the noise vector x and, therefore, has failed to learn the variability in the training data set 1.

2. Training Instability: One of the primary problems that researchers have describing their work with GANs is that the process that has to be followed to train the networks is not easy. It still fluctuates during the training process and sometimes results in positive values for difference in loss function.

3. Vanishing Gradients: However, in the training process, the discriminator can be trained to be very effective, and this leads to one major training problem; the vanishing gradients problem of the generator that hinders further improvement. This is because the discriminator will trust its own abilities and thus the gradients passed back to the generator will be tiny.

4. Non-Convergence: This gives rise to what in the context of GANs is referred to as the unstable state I, in which the two networks will not reach optima.

5. Evaluation Difficulties: The assessment of GANs is not simple, and there are several challenges that are associated with it. In contrast with most other machine learning models of text generation, there is no straightforward objective measure of the 'perceived quality' of the synthetic samples that can be easily defined in terms of a loss function. This makes it difficult to assess the performance of GANs based on certain criteria and compare them with other Convolutional GAN models.

6. High Computational Cost: Training and using GANs may require a significant amount of computation frequently during the training period. This is because, in order to train two

spiking neural networks simultaneously, the problem needs to be solved on relatively large datasets, which is rather demanding in terms of computational resources.

7. Ethical Concerns: Regarding the applications of GANs, it is the possibility provided by the networks' ability to generate realistic data for the development of such ethical concerns as deepfakes, which can be used to spread fake news and different kinds of vulnerability.

7. FUTURE DIRECTION

Generative Adversarial Networks or GANs are one of the rapidly developing fields as much effort is being given on the stability and robustness of GANs. Ongoing experiments with a variety of architectures, loss functions and even regularization to solve problems such as mode collapse and non-convergence. Conditional GANs are also being worked upon as the focus in these is on the manipulation of specific attributes in generated data which is the core of attribute manipulation for images or attribute-specific text-to-image synthesis.

Some of the most important elements of GANs are Ethical issues as their use shall be more profound in the future. There are attempts to think of ways by which fake deepfakes that look exactly like the real people cannot be created and real efforts that need to be taken to remove bias from the training data sets. Exploring cross-modal and multimodal GANs has paved way for complex multimedia content synthesis as these networks are now the source of understanding and transforming from one data type to another, text, image, and audio.

The studies that combine GANs with federated learning point to further development in the context of privacy since federated learning enables GANs to learn from multiple sources without compromising the users' privacy. This shift towards practical applications is trending, making GANs go beyond research laboratories to work in sectors such as the medical field for drug development and diagnostics, and in city planning for modeling the environment and cities.

Furthermore, the fusion of GANs with other AI technologies such as reinforcement learning and transfer learning is anticipated to unlock new capabilities and boost the efficacy of generative models. Hierarchical GANs, which employ stacked architectures, are poised to facilitate the creation of intricate and high-dimensional outputs, including lifelike 3D objects and high-fidelity videos, marking a significant leap forward in the field.

8. CONCLUSION

In the complex field of artificial intelligence especially machine learning, the Generative Adversarial Networks (GANs) have emerged as an immense invention. This methodical research paper has looked at the design and mathematical equations of these models and the versatility of the models explored in this study but has also placed them in the spirit of the contemporary research domain where GANs are being incorporated with the latest deep learning platforms and are branching out operating in areas other than image synthesis. The future of GANs is rather bright, as it can contribute to such fields as data privacy, medical

imaging, digital arts, and will be likely to help combat the problem of deepfakes. It only sweetens the technological outlook of the pillar as generative models come with the future of advanced creativity and intelligence in machines.

REFERENCES

- [1] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, Yoshua Bengio, Generative Adversarial Networks, 10 June 2014.
- [2] Ankan Dash, Junyi Ye, Guiling Wang, A review of Generative Adversarial Networks (GANs) and its applications in a wide variety of disciplines -- From Medical to Remote Sensing, 1 Oct, 2021.
- [3] Tanujit Chakraborty, Ujjwal Reddy K S, Shraddha M. Naik, Madhurima Panja, and Bayapureddy Manvitha, Ten Years of Generative Adversarial Nets (GANs): A survey of the state-of-the-art, Aug, 2023.
- [4] Zhengwei Wang, Qi She, Tom´ as E. Ward, Generative Adversarial Networks in Computer Vision: A Survey and Taxonomy, 29 Dec 2020.
- [5] Sukarna Barua, Sarah Monazam Erfani, James Bailey, FCC-GAN: A Fully Connected and Convolutional Net Architecture for GANs, 27 May 2019.
- [6] Monireh Mohebbi Moghadam, Bahar Boroomand, Mohammad Jalali, Arman Zareian, Alireza Daeijavad, and Mohammad Hossein Manshaei, Game of GANs: Game-Theoretical Models for Generative Adversarial Networks, 15 June 2021.
- [7] Monireh Mohebb, Moghadam, Bahar Boroomand, Mohammad Jalali, Arman Zareian, Alireza Daeijavad, Mohammad Hossein Manshaei, and Marwan Krunch, Game of GANs: Game-Theoretical Models for Generative Adversarial Networks, 3 Jan 2022.
- [8] Shaojie Li, Jie Wu, Xuefeng Xiao, Fei Chao, Xudong Mao, Rongrong Ji, Revisiting Discriminator in GAN Compression: A Generator-discriminator Cooperative Compression Scheme, 27 Oct 2021.
- [9] Ceyuan Yang, Yujun Shen, Yinghao Xu, Deli Zhao, BoDai, Bolei Zhou, Improving GANs with A Dynamic Discriminator, 20 Sep 2022.
- [10] Guillermo Iglesias, Edgar Talavera, A survey on GANs for computer vision: Recent research, analysis and taxonomy, 16 Feb, 2024.