Generative Advesarial Network

Tushar B. Kute,
http://tusharkute.com

Generative Adversarial Network

- Generative Adversarial Networks (GANs) are a powerful class of neural networks that are used for unsupervised learning.
- It was developed and introduced by Ian J. Goodfellow in 2014.
- GANs are basically made up of a system of two competing neural network models which compete with each other and are able to analyze, capture and copy the variations within a dataset.

Generative Adversarial Network

- It has been noticed most of the mainstream neural nets can be easily fooled into misclassifying things by adding only a small amount of noise into the original data.
- Surprisingly, the model after adding noise has higher confidence in the wrong prediction than when it predicted correctly.
- The reason for such adversary is that most machine learning models learn from a limited amount of data, which is a huge drawback, as it is prone to overfitting.
- Also, the mapping between the input and the output is almost linear. Although, it may seem that the boundaries of separation between the various classes are linear, but in reality, they are composed of linearities and even a small change in a point in the feature space might lead to misclassification of data.

Generative Adversarial Network

- Generative Adversarial Networks (GANs) can be broken down into three parts:
  - Generative: To learn a generative model, which describes how data is generated in terms of a probabilistic model.
  - Adversarial: The training of a model is done in an adversarial setting.
  - Networks: Use deep neural networks as the artificial intelligence (AI) algorithms for training purpose.

Generative Adversarial Network

- Shown below is an example of a GAN. There is a database that has real 100 rupee notes. The generator neural network generates fake 100 rupee notes. The discriminator network will help identify the real and fake notes.

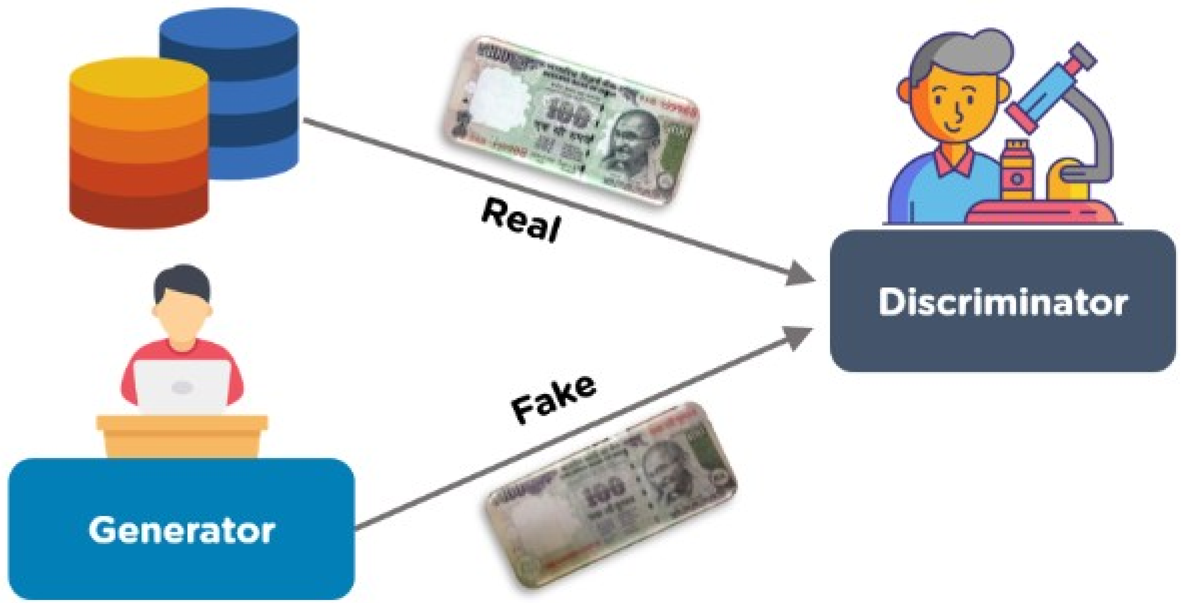

Generator

- A Generator in GANs is a neural network that creates fake data to be trained on the discriminator. It learns to generate plausible data.
- The generated examples/instances become negative training examples for the discriminator. It takes a fixed-length random vector carrying noise as input and generates a sample.

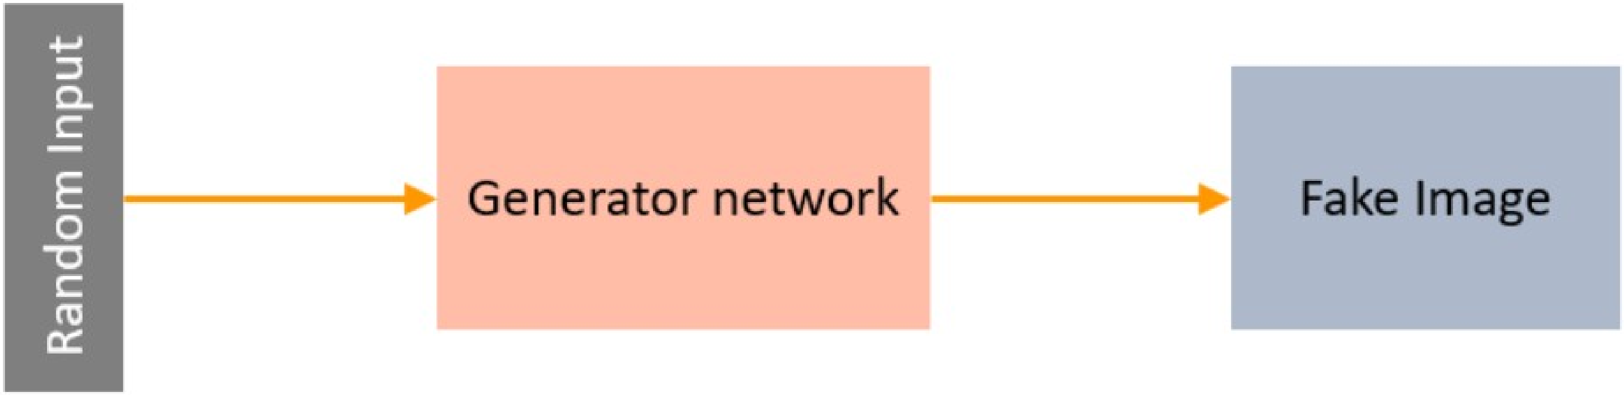

Generator

- The main aim of the Generator is to make the discriminator classify its output as real. The part of the GAN that trains the Generator includes:
  - noisy input vector
  - generator network, which transforms the random input into a data instance
  - discriminator network, which classifies the generated data
  - generator loss, which penalizes the Generator for failing to dolt the discriminator

Generator

- The backpropagation method is used to adjust each weight in the right direction by calculating the weight's impact on the output.
- It is also used to obtain gradients and these gradients can help change the generator weights.

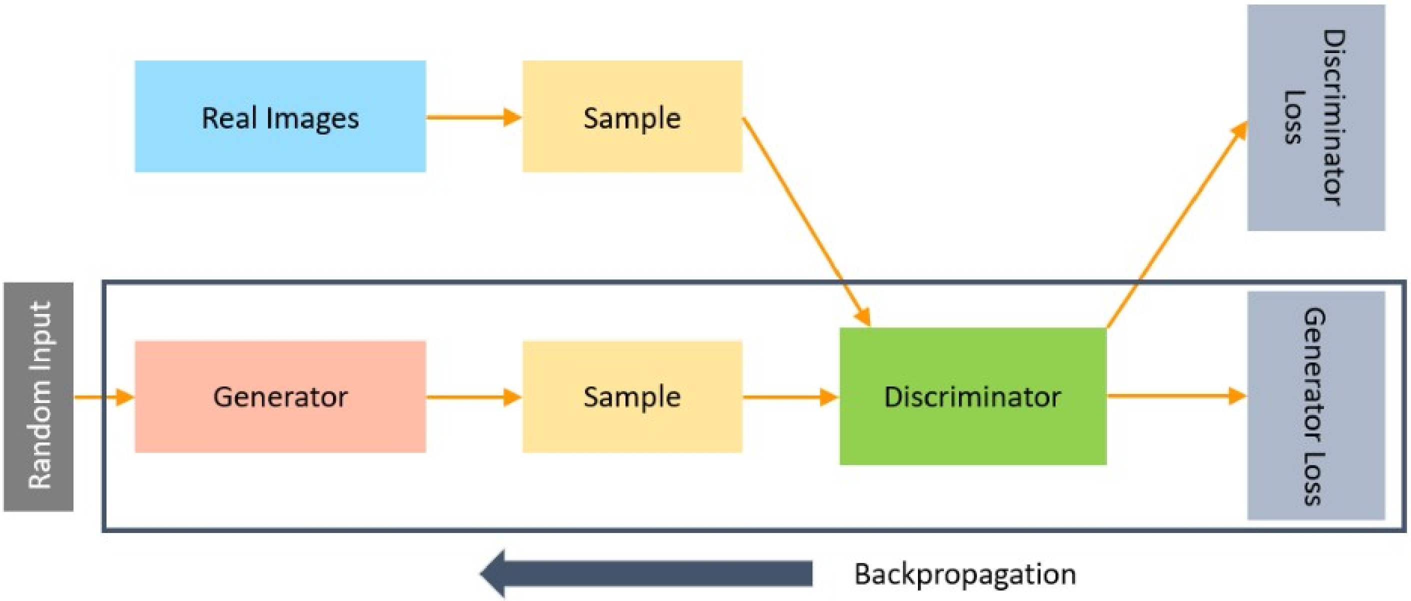

Discriminator

- The Discriminator is a neural network that identifies real data from the fake data created by the Generator. The discriminator's training data comes from different two sources:
  - The real data instances, such as real pictures of birds, humans, currency notes, etc., are used by the Discriminator as positive samples during training.
  - The fake data instances created by the Generator are used as negative examples during the training process.

Discriminator

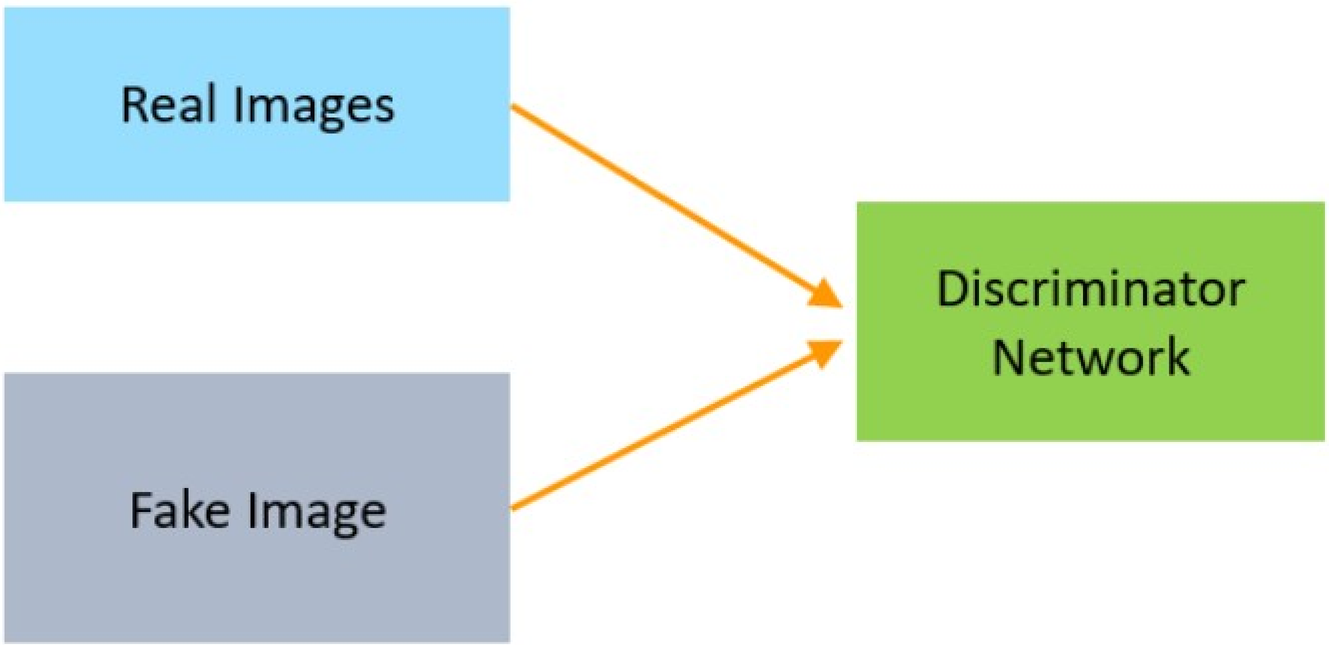

Discriminator

- While training the discriminator, it connects to two loss functions. During discriminator training, the discriminator ignores the generator loss and just uses the discriminator loss.
- In the process of training the discriminator, the discriminator classifies both real data and fake data from the generator. The discriminator loss penalizes the discriminator for misclassifying a real data instance as fake or a fake data instance as real.
- The discriminator updates its weights through backpropagation from the discriminator loss through the discriminator network.

Discriminator

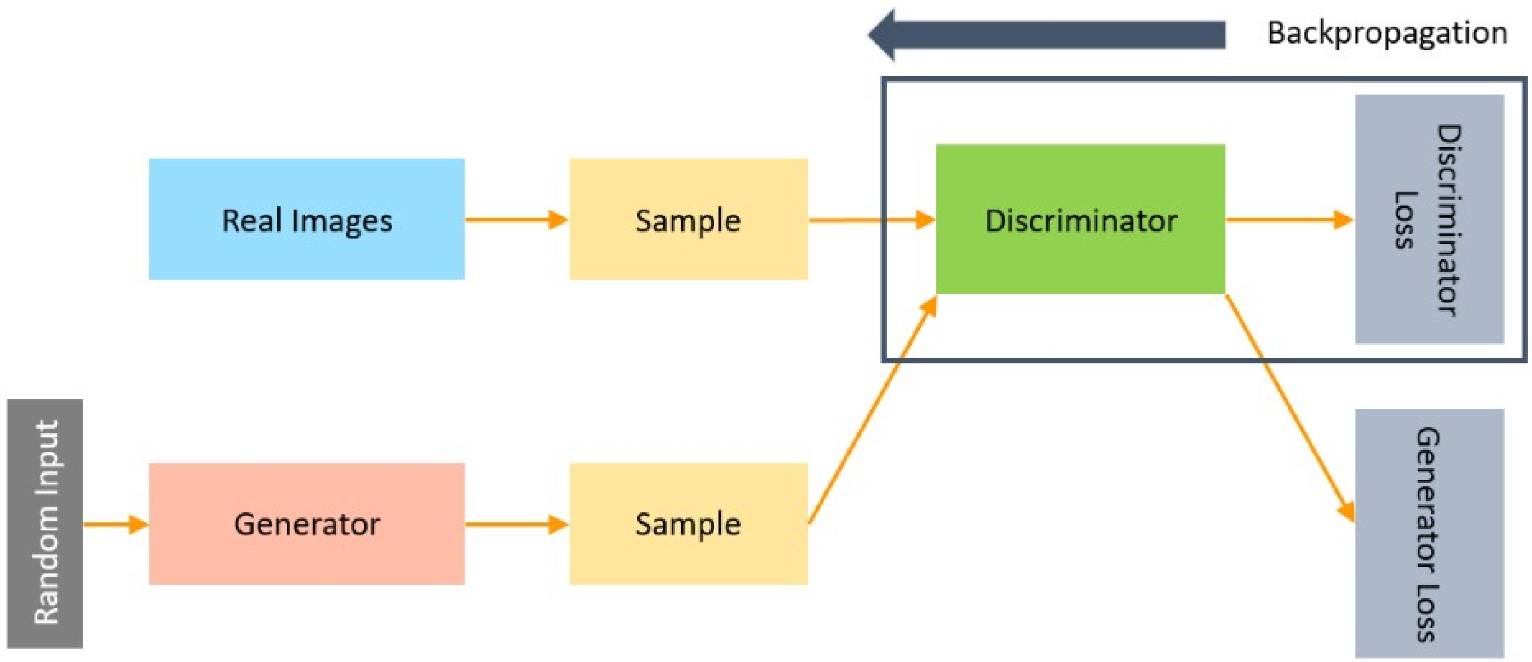

How Do GANs Work?

- GANs consists of two neural networks. There is a Generator G(x) and a Discriminator D(x).
- Both of them play an adversarial game. The generator's aim is to fool the discriminator by producing data that are similar to those in the training set.
- The discriminator will try not to be fooled by identifying fake data from real data. Both of them work simultaneously to learn and train complex data like audio, video, or image files.

How Do GANs Work?

- The Generator network takes a sample and generates a fake sample of data.
- The Generator is trained to increase the Discriminator network's probability of making mistakes.

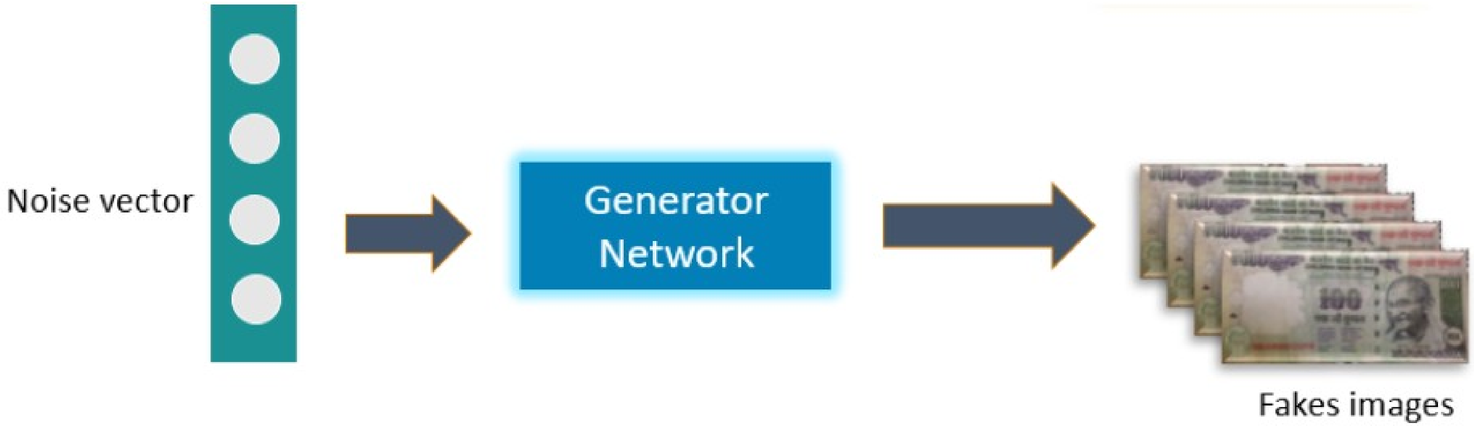

How Do GANs Work?

- Below is an example of a GAN trying to identify if the 100 rupee notes are real or fake. So, first, a noise vector or the input vector is fed to the Generator network.
- The generator creates fake 100 rupee notes. The real images of 100 rupee notes stored in a database are passed to the discriminator along with the fake notes.
- The Discriminator then identifies the notes as classifying them as real or fake.
- We train the model, calculate the loss function at the end of the discriminator network, and backpropagate the loss into both discriminator and generator models.

How Do GANs Work?

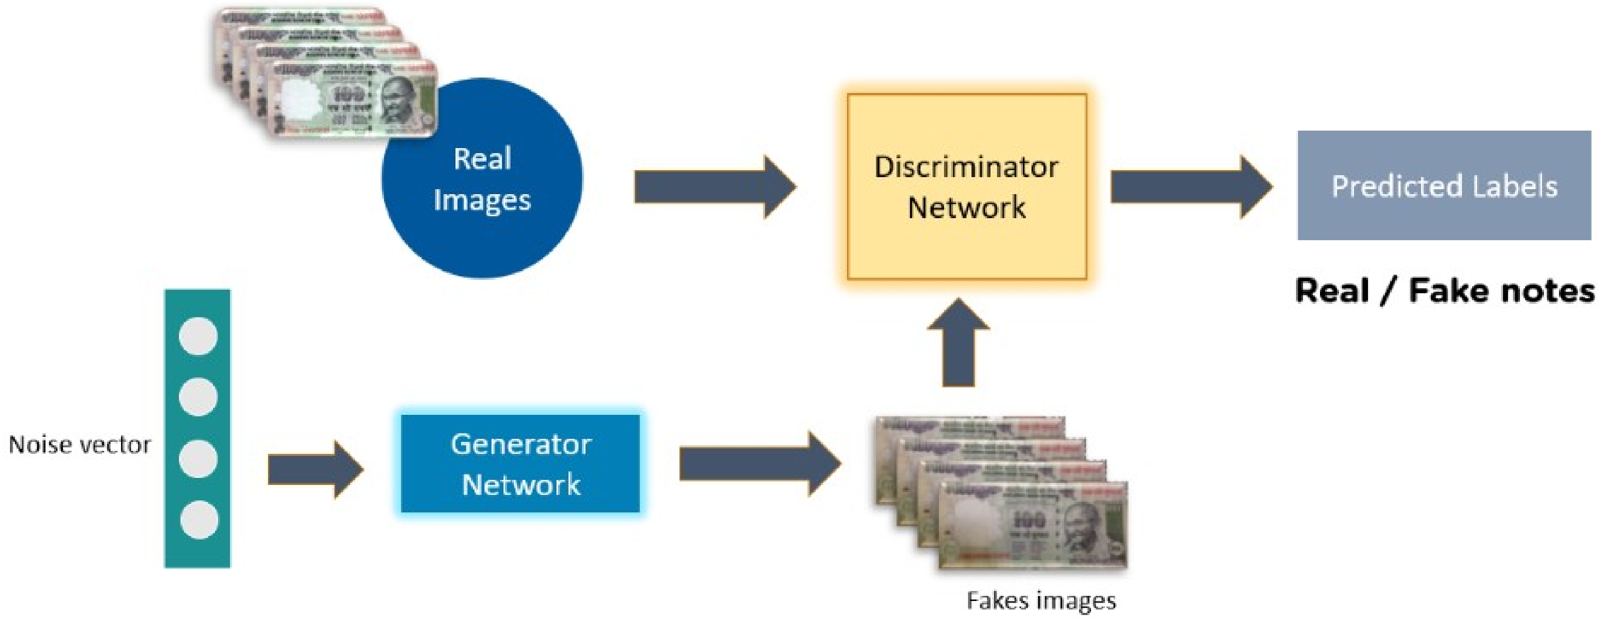

Steps for Training GAN

- Define the problem
- Choose the architecture of GAN
- Train discriminator on real data
- Generate fake inputs for the generator
- Train discriminator on fake data
- Train generator with the output of the discriminator

Types of GANs

- Vanilla GANs:
  - Vanilla GANs have a min-max optimization formulation where the Discriminator is a binary classifier and uses sigmoid cross-entropy loss during optimization.
  - The Generator and the Discriminator in Vanilla GANs are multi-layer perceptrons.
  - The algorithm tries to optimize the mathematical equation using stochastic gradient descent.

Types of GANs

- Deep Convolutional GANs (DCGANs):
  - DCGANs support convolution neural networks instead of vanilla neural networks at both Discriminator and Generator.
  - They are more stable and generate better quality images. The Generator is a set of convolution layers with fractional-strided convolutions or transpose convolutions, so it up-samples the input image at every convolutional layer.
  - The discriminator is a set of convolution layers with strided convolutions, so it down-samples the input image at every convolution layer.

Types of GANs

- Conditional GANs:
  - Vanilla GANs can be extended into Conditional models by using extra-label information to generate better results.
  - In CGAN, an additional parameter ‘y’ is added to the Generator for generating the corresponding data.
  - Labels are fed as input to the Discriminator to help distinguish the real data from the fake generated data.

Types of GANs

- Super Resolution GANs:
  - SRGANs use deep neural networks along with an adversarial network to produce higher resolution images.
  - SRGANs generate a photorealistic high-resolution image when given a low-resolution image.

GAN Applications

- With the help of DCGANs, you can train images of cartoon characters for generating faces of anime characters as well as Pokemon characters.

GAN Applications

- GANs can be trained on the images of humans to generate realistic faces. The faces that you see below have been generated using GANs and do not exist in reality.

GAN Applications

- GANs can build realistic images from textual descriptions of objects like birds, humans, and other animals.
- We input a sentence and generate multiple images fitting the description.

Thank you

@mitu_skillologies

/mITuSkillologies

@mitu_group

/company/mitu-skillologies

MITUSkillologies

Web Resources
http://mitu.co.in
http://tusharkute.com

contact@mitu.co.in

tushar@tusharkute.com