

Image-to-Image Translation using Generative Adversarial Networks (GANs)

Anant Veer Bagrodia

Vasant Valley School, Vasant Kunj, New Delhi
E-mail: anantveerbagrodia@vasantvalley.edu.in

Abstract—This paper aims to provide an overview of Generative Adversarial Networks (GANs) and a detailed explanation of their application in Image-to-Image Translation tasks.

The paper delves into the working, structure, functioning, and training process of GANs. Additionally, it includes sections on the various types of GANs, as well as their advantages and disadvantages. Exciting applications and potential challenges of GANs are also discussed.

Furthermore, an in-depth example of Image-to-Image translation, one of GANs' applications, is presented. Using Python, a GAN was developed and trained on the Fashion MNIST dataset, consisting of 70,000 grayscale images across 10 different fashion categories.

The paper explains the process of data selection and preprocessing, the architecture of the GAN, the training and evaluation process, and the obtained results and analysis. The model achieved a satisfactory discriminator loss of 0.48 and demonstrated an excellent recall score of 0.9694.

Although there is room for improvement in terms of accuracy and precision, the model exhibits potential for practical applications in the fashion industry, such as virtual try-on and fashion recommendation systems. Finally, the paper explores various techniques that can be implemented to enhance the model's performance.

WHAT ARE GENERATIVE ADVERSARIAL NETWORKS (GANS)?

Developed in 2014 by Ian Good fellow, GANs are a form of generative modeling that aim to generate a new, synthetic set of data which resembles the known data on which it was trained. GANs can be broken into three parts:

- (i) Generative - GANs describe how data is generated visually in terms of a probabilistic model.
- (ii) Adversarial - The training of the model is done in an adversarial environment (Generator and discriminator compete with each other). Game theory can also be used to analyze the working of GANs, hence they are adversarial.
- (iii) Network - GANs use deep neural networks as algorithms for training purposes.

Used for unsupervised learning, GANs systems consist of two main components (neural networks models), which compete with each other to capture, copy, and analyze the variations in a dataset. These two components are known as Generator and Discriminator. The generator utilizes random input, usually noise, to produce samples like images, audio, or text that are similar to the training data it was trained on. The generator aims to generate samples that are indistinguishable from real data for the discriminator. Thus, a generator is “a model that is used to generate new plausible examples from the problem domain”.

Conversely, the discriminator's goal is to differentiate between real samples from the domain and samples generated by the generator. It is trained with real samples from the training data, as well as generated samples produced by the generator. The training of the discriminator is a supervised approach. The primary aim of the discriminator is to accurately classify real data from the domain as real and data generated by the generator as fake. Thus, a discriminator is a model employed to categorize data samples as real (from the domain) or fake (produced by the generator).

The training process involves an adversarial game between the generator and the discriminator. The generator aims to produce samples that closely resemble real data from the domain and fool the discriminator, while the discriminator aims to improve its ability to distinguish between real and fake data. This adversarial training leads to both networks learning and improving over time.

As training progresses, the generator improves its ability of consistently generating realistic samples that fool the discriminator, while the discriminator becomes more skilled at differentiating between real and generated data. This process continues till when

the generator is capable of consistently generating high-quality samples that are difficult for the discriminator to distinguish from real data.

The generative model is trained such that it tries to maximize the probability of the Discriminator making a mistake. On the other hand, the discriminator is based on a model that “estimates the probability that the sample that it received is from the domain and not from the generator”. The GANs are formulated as a “zero-sum minimax game, where the Discriminator is trying to minimize its reward $V(D, G)$ and the Generator is trying to minimize the Discriminator’s reward/maximize its loss”.

It can be mathematically described by the formula below:

$$\min_G \max_D V(D, G)$$

$$V(D, G) = \mathbb{E}_{x \sim p_{data}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))]$$

where,

G = Generator

D = Discriminator

Pdata(x) = distribution of real data

P(z) = distribution of generator

x = sample from Pdata(x)

z = sample from P(z)

D(x) = Discriminator network

G(z) = Generator network

In this case, zero-sum means that when the discriminator successfully identifies real and fake samples, it is rewarded (no change is needed to the model parameters), whereas the generator is penalized (large updates to model parameters).

Conversely, when the generator fools the discriminator, it is rewarded, but the discriminator is penalized.

Different Types of GAN Models

- Vanilla GAN: Vanilla GAN is the simplest form of GAN. The algorithm simply tries to optimize the mathematical equation using stochastic gradient descent.
- Conditional GAN (CGAN): CGAN is a GAN in which some additional conditional parameters are present. In CGAN, an additional parameter ‘y’ is added to the Generator for producing the corresponding data. Additionally, labels are added to the input to the Discriminator to help it distinguish the real data from the fake generated data.
- Deep Convolutional GAN (DCGAN): DCGAN is one of the most popular and successful implementations of GAN. Instead of multi-layer perceptrons, it consists of ConvNets. The layers are not fully connected.
- Laplacian Pyramid GAN (LAPGAN): The Laplacian pyramid is a “linear invertible image representation consisting of a set of band-pass images, spaced an octave apart, plus a low-frequency residual”. The main benefit of this approach is that it produces very high-quality images. First, the image is down-sampled at each layer of the pyramid. Then, in a backward pass, it is again up-scaled at each layer. At these layers, the image acquires some noise from the Conditional GAN until it reaches back to its original size.
- Super Resolution GAN (SRGAN): In addition to an adversarial network, SRGAN uses a deep neural network along with it in order to produce higher-resolution images. Its primary application is in the up-scaling of low-resolution images to enhance their details while also minimizing errors in doing so. This is known as domain transformation.
- Least Square GAN(LSGAN) – LSGAN adopts the least-square loss function for the discriminator.
- Auxiliary Classifier GAN(ACGAN) – It is an advanced version of CGAN. In ACGAN, the Discriminator functions not only to classify the image as real/fake, but also to provide input image’s source/class label.
- Dual Video Discriminator GAN – DVD-GAN is utilized mainly for high-definition video generation. It is implemented using the BigGAN architecture. It consists of two discriminators: a Spatial Discriminator and a Temporal Discriminator.

ADVANTAGES OF GENERATIVE ADVERSARIAL NETWORKS (GANS):

- Generation of synthetic data: GANs generate a new, synthetic set of data which resembles the known data on which it was trained. This is useful for data augmentation, anomaly detection, or creative applications.
- Producing high quality results: GANs are capable of consistently producing high-quality and photorealistic results in tasks such as image synthesis, video synthesis, music synthesis, etc.
- Applications in unsupervised learning: Since GANs can be trained using only unlabeled data, they can be utilized for unsupervised learning tasks, where labeled data is not available for use.

- Versatility and scope: GANs have a huge number of vast applications, such as image and video synthesis, anomaly detection, data augmentation, image-to-image translation, etc.

DISADVANTAGES OF GENERATIVE ADVERSARIAL NETWORKS (GANS):

- Training instability: Many difficulties can arise in the training of GANs, such as the risk of instability, mode collapse, or failure to converge.
- High computational cost: When generating high-resolution images or operating on large datasets, GANs can require a lot of computational resources and take a long time to train.
- Overfitting: GANs may overfit the training data. This results in them producing synthetic data that is almost identical to the training data and lacking diversity.
- Bias and Unfairness: GANs can reflect the biases and unfairness present in the training data. This can result in them generating synthetic data that is discriminatory or biased.

WHY WERE GANS DEVELOPED?

By adding some amount of noise to data, machine learning algorithms and neural networks can easily be fooled to misclassify things and the chances of this happening increase.

In complex domains or domains with a limited amount of data, generative modeling allows for access to more training data for modeling. GANs have seen much success in this use case in domains such as deep reinforcement learning.

There are numerous reasons why GANs are significant and require further study. The creator of GANs, Ian Goodfellow himself outlined a number of these in his 2016 conference keynote and associated technical report titled “NIPS 2016 Tutorial: Generative Adversarial Networks.”

He focussed on GANs’ ability to “model high-dimensional data, handle missing data, and provide multi-modal outputs or multiple plausible answers”.

One of the most compelling applications of GANs is conditional GANs. Conditional GANs are used for tasks that require the generation of new examples, such as creating art and image-to-image translation.

APPLICATIONS OF GENERATIVE ADVERSARIAL NETWORKS (GANS)

- Generation of new data from available data – GANs generate a new, synthetic set of data which resembles the known data on which it was trained.
- GANs are not just limited to images, they can also be applied to generate text, articles, songs, poems, etc.
- Generation of music by using a clone voice – If you input some voice to GANs, they can produce a similar clone feature of it. Researchers from NIT in Tokyo proposed a system that is able to generate melodies from lyrics with the help of learned relationships between notes and subjects.
- Text to Image Generation (Object GAN and Object Driven GAN).
- GANs are used for the creation of anime characters in game development as well as in animation production.
- Image to Image Translation – GANs can be used to translate one image to another without changing the background of the source image. For example, GANs can replace a tiger in a forest with a leopard in the same forest.
- If you input a low-resolution image or video, GANs can generate a high-resolution version of the same.
- By training GANs on small frames of videos, they are capable of predicting and generating the next frame of the videos.
- Recently, researchers from the College of London published a system called GAN-TTS that learns to generate raw audio through training on 567 corpora of speech data.

Research on GANs is progressing rapidly and development is advancing each day. Already, Microsoft has collaborated with OpenAI to work on GPT and explore the power of GANs.

CHALLENGES FACED BY GENERATIVE ADVERSARIAL NETWORKS (GANS)

- The problem of stability between generator and discriminator - It is challenging to find the right equilibrium, as the discriminator should be neither too strict (overpower the generator and prevent it from producing realistic images), nor too lenient (may fail to distinguish between real and fake samples, leading to poor quality of generated images).
- The problem in understanding of global objects – GANs struggle to understand the global structure or holistic features in images. As a result, GANs may generate images that are unrealistic or impossible.
- GANs have limitations in understanding and generating 3D objects. They are primarily designed to work with 1D or 2D images.

IMAGE-TO-IMAGE TRANSLATION USING GANS

Image-to-image translation is a task that involves “transforming an input image from one domain to another domain while still maintaining the relevant details and semantic information”. It has practical applications in various fields, such as converting grayscale images to color, turning sketches into realistic images, or translating images between different artistic styles. Researchers have developed GAN architectures such as Pix2Pix, CycleGAN, and UNIT, for image translation tasks.

We will now be looking at a practical implementation of GANs for image-to-image translation using Python:

DATASET SELECTION AND PREPROCESSING

For this example, I chose the "Fashion MNIST" dataset. It consists of 70,000 grayscale images of 10 different fashion categories. (Programmed with assistance from ChatGPT)

CODE:

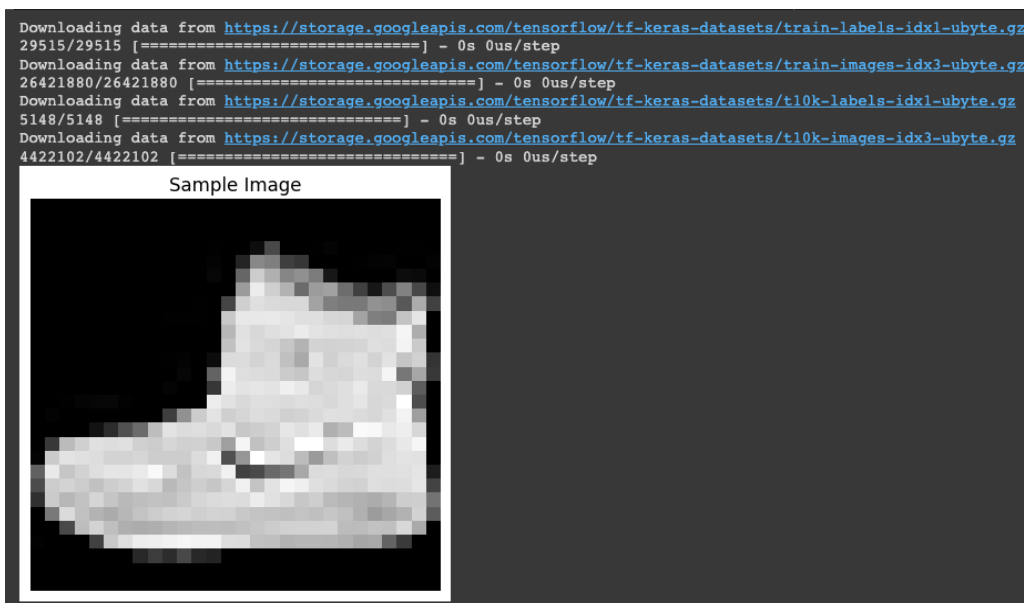
```
1 import tensorflow as tf
2 import matplotlib.pyplot as plt
3
4 # Load the Fashion MNIST dataset
5 (x_train, y_train), (x_test, y_test) = tf.keras.datasets.fashion_mnist.load_data()
6
7 # Reshape the images to (28, 28, 1) and normalize pixel values to the range [0, 1]
8 x_train = x_train.reshape(-1, 28, 28, 1) / 255.0
9 x_test = x_test.reshape(-1, 28, 28, 1) / 255.0
10
11 # Convert the labels to one-hot encoded vectors
12 y_train = tf.keras.utils.to_categorical(y_train)
13 y_test = tf.keras.utils.to_categorical(y_test)
14
15 # Visualize a sample from the dataset
16 plt.imshow(x_train[0].reshape(28, 28), cmap='gray')
17 plt.title('Sample Image')
18 plt.axis('off')
19 plt.show()
20
```

In the above code, firstly the “tf.keras.datasets.fashion_mnist.load_data()” function is used to load the Fashion MNIST dataset. The function also splits the dataset into training and testing sets. To match the input shape expected by most convolutional neural networks, the images are reshaped to (28, 28, 1) and the pixel values are normalized to the range [0, 1].

Using the tf.keras.utils.to_categorical() function, the labels are converted to one-hot encoded vectors.

Lastly, using plt.imshow(), a random sample image from the dataset is visualized.

OUTPUT:



GAN ARCHITECTURE

CODE:

```

1 import tensorflow as tf
2 from tensorflow.keras import layers
3
4 # Generator model
5 generator = tf.keras.Sequential()
6 generator.add(layers.Dense(256, input_shape=(100,), activation='relu'))
7 generator.add(layers.Dense(784, activation='tanh'))
8 generator.add(layers.Reshape((28, 28, 1)))
9
10 # Discriminator model
11 discriminator = tf.keras.Sequential()
12 discriminator.add(layers.Flatten(input_shape=(28, 28, 1)))
13 discriminator.add(layers.Dense(256, activation='relu'))
14 discriminator.add(layers.Dense(1, activation='sigmoid'))
15
16 # Combined GAN model
17 gan = tf.keras.Sequential([generator, discriminator])
18
19 # Compile the discriminator
20 discriminator.compile(optimizer=tf.keras.optimizers.Adam(0.0002, 0.5),
21                       loss='binary_crossentropy',
22                       metrics=['accuracy'])
23
24 # Compile the GAN
25 gan.compile(optimizer=tf.keras.optimizers.Adam(0.0002, 0.5),
26             loss='binary_crossentropy')
27
28 # Print the model summaries
29 generator.summary()
30 discriminator.summary()
31 gan.summary()
32

```

In the above code, we define our GAN architecture consisting of a generator and a discriminator. The generator is responsible for generating plausible grayscale images of fashion items from random noise, while the discriminator aims to distinguish between real and generated fashion items. Both the generator and discriminator are composed of several layers, including convolutional and fully connected layers.

The generator model takes random noise as input, passes it through dense layers with activation functions, reshapes the generated image, and outputs the final generated image. The discriminator model takes an input image, flattens it, passes it through dense layers, and outputs a binary classification result indicating whether the image is real or fake.

The GAN model combines the generator and discriminator, where the generator tries to fool the discriminator by generating realistic images, and the discriminator tries to accurately classify the real and fake images.

The discriminator and GAN models are compiled with appropriate optimizers and loss functions. Lastly, the generator, discriminator, and GAN model summaries are printed.

OUTPUT:

```

Model: "sequential"
-----
Layer (type)                Output Shape                Param #
-----
dense (Dense)                (None, 256)                 25856
dense_1 (Dense)              (None, 784)                 201488
reshape (Reshape)           (None, 28, 28, 1)          0
-----
Total params: 227,344
Trainable params: 227,344
Non-trainable params: 0
-----
Model: "sequential_1"
-----
Layer (type)                Output Shape                Param #
-----
flatten (Flatten)           (None, 784)                 0
dense_2 (Dense)              (None, 256)                 200960
dense_3 (Dense)              (None, 1)                   257
-----
Total params: 201,217
Trainable params: 201,217
Non-trainable params: 0
-----
Model: "sequential_2"
-----
Layer (type)                Output Shape                Param #
-----
sequential (Sequential)     (None, 28, 28, 1)          227344
sequential_1 (Sequential)   (None, 1)                   201217
-----
Total params: 428,561
Trainable params: 428,561
Non-trainable params: 0
-----

```

MODEL TRAINING AND EVALUATION:

I trained the GAN model for 50 epochs (complete pass through the entire training dataset) with a batch size (number of samples processed in each iteration) of 128. Even though a higher number of epochs would likely make the model more accurate, I chose 50 as the number of epochs because I felt it is the right tradeoff between execution time (1 hour, 5 minutes) and model performance.

```

1 import numpy as np
2 # Set the number of epochs and batch size
3 epochs = 50
4 batch_size = 128
5 latent_dim = 100
6 # Training loop
7 for epoch in range(epochs):
8     for batch in range(len(x_train) // batch_size):
9         # Train the discriminator
10        noise = tf.random.normal(shape=(batch_size, latent_dim))
11        fake_images = generator.predict(noise)
12
13        real_labels = tf.ones((batch_size, 1))
14        fake_labels = tf.zeros((batch_size, 1))
15
16        d_loss_real = discriminator.train_on_batch(x_train[batch * batch_size : (batch + 1) * batch_size], real_labels)
17        d_loss_fake = discriminator.train_on_batch(fake_images, fake_labels)
18        discriminator_loss = 0.5 * np.add(d_loss_real, d_loss_fake)
19
20        # Train the generator
21        noise = tf.random.normal(shape=(batch_size, latent_dim))
22        labels = tf.ones((batch_size, 1))
23        generator_loss = gan.train_on_batch(noise, labels)
24
25        # Print the loss for each epoch
26        print(f"Epoch {epoch+1}/{epochs} - Discriminator Loss: {discriminator_loss}, Generator Loss: {generator_loss}")
27
28 # Evaluate the discriminator
29 test_noise = tf.random.normal(shape=(len(x_test), latent_dim))
30 fake_images = generator.predict(test_noise)
31 real_labels = tf.ones((len(x_test), 1))
32 fake_labels = tf.zeros((len(x_test), 1))
33
34 d_loss_real = discriminator.evaluate(x_test, real_labels)
35 d_loss_fake = discriminator.evaluate(fake_images, fake_labels)
36 discriminator_loss = 0.5 * np.add(d_loss_real, d_loss_fake)
37
38 # Print the evaluation metrics
39 print("Discriminator Loss:", discriminator_loss)
40

```

In the above code, we use a loop to train the GAN model to generate realistic fashion items based on the input images. The discriminator is trained on both real and fake images, and the generator is trained to generate more realistic images.

During training, the generator and discriminator are updated alternately. For each epoch, the training loop iterates over the batches in the training dataset and in each batch, the discriminator is trained first.

The discriminator is trained using a combination of real images from the training dataset and fake images produced by the generator. Random noise is used as input to the generator to generate fake images. The discriminator is then trained on batches of real images with labels set to 1 (`real_labels`) and fake images with labels set to 0 (`fake_labels`).

The discriminator loss is calculated as the average of the losses obtained from training on real and fake samples. It measures the discriminator's ability to distinguish between real and generated images.

The discriminator aimed to minimize this loss by correctly classifying real and generated fashion items. On the other hand, the generator sought to maximize the discriminator's loss, encouraging the generation of more realistic fashion items.

Next, the generator is trained. Random noise of shape (`batch_size`, `latent_dim`) is used as input to the generator. The labels for the generator are set to 1 (`labels`), indicating that the generated images should be classified as real. The generator loss is computed using the `train_on_batch` method of the GAN model, which updates the generator's weights based on its performance in fooling the discriminator.

After each epoch, the discriminator loss and generator loss are printed to track the progress of the training.

Once the training is complete, the discriminator's performance is evaluated on the test dataset. Random noise of shape $(\text{len}(x_test), \text{latent_dim})$ is generated and used as input to the generator to produce fake images. The discriminator is then evaluated on both the real test images and the generated fake images.

The discriminator loss is computed by calculating the loss separately for the real test images and the fake images. They are then averaged using the `np.add` function and multiplied by 0.5 to get the final discriminator loss.

Lastly, the discriminator loss is printed as the evaluation metric for assessing the discriminator's ability to distinguish between real and generated images.

```
Epoch 50/50 - Discriminator Loss: [0.85524076 0.4765625 ], Generator Loss: 1.1077752113342285
313/313 [=====] - 1s 3ms/step
313/313 [=====] - 1s 3ms/step - loss: 0.1743 - accuracy: 0.9694
313/313 [=====] - 1s 3ms/step - loss: 1.9389 - accuracy: 0.0075
Discriminator Loss: [1.05659515 0.48844999]
```

We got a final value of 0.48 as the Discriminator Loss, implying that our model demonstrates promising discriminative capabilities (a lower value suggests better discrimination capabilities).

RESULTS AND ANALYSIS

Code:

```
1 from sklearn.metrics import accuracy_score, precision_score, f1_score, recall_score
2
3 # Generate a batch of new images using the trained generator
4 num_samples = 10
5 noise = tf.random.normal(shape=(num_samples, latent_dim))
6 generated_images = generator.predict(noise)
7
8 # Display the generated images
9 fig, axs = plt.subplots(1, num_samples, figsize=(10, 2))
10 for i in range(num_samples):
11     axs[i].imshow(generated_images[i, :, :, 0], cmap='gray')
12     axs[i].axis('off')
13 plt.show()
14
15 # Evaluate the discriminator on the test set
16 test_noise = tf.random.normal(shape=(len(x_test), latent_dim))
17 fake_images = generator.predict(test_noise)
18 real_labels = tf.ones((len(x_test), 1))
19 fake_labels = tf.zeros((len(x_test), 1))
20
21 # Calculate the discriminator predictions
22 real_predictions = discriminator.predict(x_test)
23 fake_predictions = discriminator.predict(fake_images)
24
25 # Compute evaluation metrics
26 real_predictions = np.round(real_predictions)
27 fake_predictions = np.round(fake_predictions)
28
29 accuracy = accuracy_score(np.concatenate((np.ones(len(x_test)), np.zeros(len(fake_images)))), np.concatenate((real_predictions, fake_predictions)))
30 precision = precision_score(np.concatenate((np.ones(len(x_test)), np.zeros(len(fake_images)))), np.concatenate((real_predictions, fake_predictions)))
31 recall = recall_score(np.concatenate((np.ones(len(x_test)), np.zeros(len(fake_images)))), np.concatenate((real_predictions, fake_predictions)))
32 f1 = f1_score(np.concatenate((np.ones(len(x_test)), np.zeros(len(fake_images)))), np.concatenate((real_predictions, fake_predictions)))
33
34 # Print the evaluation metrics
35 print("Accuracy:", accuracy)
36 print("Precision:", precision)
37 print("Recall:", recall)
38 print("F1-Score:", f1)
39
```

Output:

The evaluation results for our model are as follows:

Accuracy: 0.48735

Precision: 0.49355939106970115

Recall: 0.9694

F1-Score: 0.6540939914307884

The high recall value (also known as sensitivity or true positive rate), proportion of actual positive instances correctly predicted by the model, of 0.9694 suggests that our model excelled in correctly identifying positive instances (fashion items) from the dataset. This implies that the model has a low rate of false negatives and is successful in minimizing the instances where positive samples are missed.

However, despite the high recall, our model exhibited lower accuracy (overall correctness of the model's predictions) and precision (proportion of correctly predicted positive instances out of the total instances predicted as positive) values. The accuracy value of 0.48735 indicates that the overall correctness of the model's predictions was relatively low, achieving accurate classifications for only approximately 48.735% of the samples in the test dataset. Additionally, the precision value of 0.49355939106970115 highlights that when the model predicted an instance as positive (a fashion item), only around 49.36% of these predictions were correct.

The discrepancy between the high recall and the lower accuracy and precision can be attributed to a higher number of false positives. This means that the model may have mistakenly classified negative instances (non-fashion items) as positive instances (fashion items). This situation is prevalent in imbalanced datasets or cases where the cost of false negatives is considered to be higher than false positives.

The F1-score, which provides a balanced measure of precision and recall, was calculated to be 0.6540939914307884. This suggests a moderate balance between precision and recall, indicating that the model achieved reasonable performance in capturing both true positive and true negative instances.

To improve the model's performance, further analysis and refinement is necessary. Techniques such as adjusting the classification threshold, incorporating class weights, and adjusting the number of epochs could help in optimizing the trade-off between accuracy, precision, and recall.

We also need to analyze the misclassified instances, explore feature engineering, or consider using more advanced architectures or pre-trained models.

CONCLUSION

In conclusion, the GAN model trained on the Fashion MNIST dataset achieved a satisfactory discriminator loss of 0.48. The model also demonstrated an excellent recall score of 0.9694. While there is room for improvement in terms of accuracy and precision, the model shows potential for applications in the fashion industry, such as virtual try-on and fashion recommendation systems. Further refinements and optimizations can enhance its performance and make it suitable for practical use. To improve the model's performance, techniques such as adjusting the classification threshold, incorporating class weights, increasing the number of epochs, or considering more advanced architectures or pre-trained models can be explored. Additionally, analyzing misclassified instances, performing feature engineering, or increasing the model's complexity could lead to enhanced performance.

REFERENCES

- [1]. <https://www.geeksforgeeks.org/generative-adversarial-network-gan/>
- [2]. <https://machinelearningmastery.com/what-are-generative-adversarial-networks-gans/>
- [3]. <https://www.analyticsvidhya.com/blog/2021/10/an-end-to-end-introduction-to-generative-adversarial-networks-gans/>
- [4]. https://developers.google.com/machine-learning/gan/gan_structure
- [5]. <https://arxiv.org/abs/1406.2661>
- [6]. <https://www.simplilearn.com/tutorials/deep-learning-tutorial/generative-adversarial-networks-gans>
- [7]. <https://www.techtarget.com/searchenterpriseai/definition/generative-adversarial-network-GAN>
- [8]. <http://papers.neurips.cc/paper/5423-generative-adversarial-nets.pdf>
- [9]. <https://towardsdatascience.com/understanding-generative-adversarial-networks-gans-cd6e4651a29>
- [10]. <https://www.javatpoint.com/generative-adversarial-network>
- [11]. <https://wiki.pathmind.com/generative-adversarial-network-gan>
- [12]. <https://medium.com/analytics-vidhya/gans-a-brief-introduction-to-generative-adversarial-networks-f06216c7200e>
- [13]. <https://www.leewayhertz.com/generative-adversarial-networks/>