



Image generation with Generative adversarial networks

Training GANs for Image Generation

Generative Adversarial Networks (GANs) consist of two neural networks—the Generator and the Discriminator—that compete with each other. The Generator creates images from random noise aiming to make them look as realistic as possible while the Discriminator evaluates these images to determine whether they are real or fake.

Training Generative Adversarial Networks (GANs) is an iterative process that revolves around the interaction between two neural networks:

Training the Discriminator

The Discriminator starts by being trained on a dataset containing real images. Its goal is to differentiate between these real images and fake images generated by the Generator. Through backpropagation and gradient descent, the Discriminator adjusts its parameters to improve its ability to accurately classify real and generated images.

Training the Generator

Concurrently the Generator is trained to produce images that are increasingly difficult for the Discriminator to distinguish from real images. Initially the Generator generates random noise but as training progresses it learns to generate images that resemble those in the training dataset. The Generator's parameters are adjusted based on the feedback from the Discriminator, optimizing the Generator's ability to create more realistic and high-quality images.

Implementing Generative Adversarial Networks (GANs) for Image Generation

Step 1: Import Necessary Libraries and Load Dataset

Import necessary libraries including TensorFlow, Keras layers and models, NumPy for numerical operations and Matplotlib for plotting.

```
import numpy as np
import tensorflow as tf
from tensorflow.keras import layers, models, optimizers
import matplotlib.pyplot as plt
```



DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING

Step 2: Dataset Preparation

Proper data preparation is crucial for the successful training of neural networks. For the MNIST dataset the preprocessing steps include loading the dataset reshaping the images to ensure they are in the correct format for TensorFlow processing and normalizing the pixel values to the range [0,1]. Normalization helps stabilize the training process by keeping the input values small.

```
(x_train, _), (_, _) = tf.keras.datasets.mnist.load_data()  
x_train = x_train.reshape((-1, 28, 28, 1)).astype('float32') / 255.0
```

Step 2: Building the Models

This step involves defining the architecture for both the generator and the discriminator using convolutional neural network (CNN) layers, tailored to efficiently process and generate image data.

Generator Model with CNN Layers

The generator's role in a GAN is to synthesize new images that mimic the distribution of a given dataset. In this case, we use convolutional transpose layers, which are effective for upscaling the input and creating detailed images from a lower-dimensional noise vector.

- Dense Layer: Converts the input 100-dimensional noise vector into a high-dimensional feature map.
- Reshape: Transforms the feature map into a 3D shape that can be processed by convolutional layers.
- Conv2DTranspose Layers: These layers perform upscaling and convolution simultaneously, gradually increasing the resolution of the generated image.
- BatchNormalization: Stabilizes the learning process and helps in faster convergence.
- Activation Functions: 'ReLU' is used for non-linearity in intermediate layers, while 'sigmoid' is used in the output layer to normalize the pixel values between 0 and 1.

```
def build_generator_cnn():  
    model = models.Sequential([  
        layers.Dense(7*7*128, input_dim=100, activation='relu'),  
        layers.Reshape((7, 7, 128)), # Reshape into an initial image format  
  
        layers.Conv2DTranspose(128, kernel_size=4, strides=2, padding='same', activation='relu'),
```



```
layers.BatchNormalization(),
```

```
layers.Conv2DTranspose(128, kernel_size=4, strides=2, padding='same', activation='relu'),
```

```
layers.BatchNormalization(),
```

```
layers.Conv2D(1, kernel_size=7, activation='sigmoid', padding='same')
```

```
)
```

```
return model
```

Discriminator Model with CNN Layers

The discriminator is a binary classifier that determines whether a given image is real (from the dataset) or fake (generated by the generator).

- Conv2D Layers: Perform convolutions with a stride of 2 to downsample the image, reducing its dimensionality and increasing the field of view of the filters.
- BatchNormalization: Used here as well to ensure stable training.
- Flatten: Converts the 2D feature maps into a 1D feature vector necessary for classification.
- Dense Output Layer: Outputs a single probability indicating the likelihood that the input image is real.