Master 2 Computer Science

Advanced Neural Networks Architectures

Generative Models

Or how magic happens in AI

Part 3. Generative Adversarial Networks

Akka Zemmari

Overview

- Introduction to GANs
- GAN Architecture
- Mathematical Formulation
- Training Process
- Applications and Use Cases
- Challenges and Future Directions

Introduction to GANs

- GANs were introduced by Ian Goodfellow et al. in 2014.
- They are a class of machine learning frameworks for generative modeling.
- Consist of two neural networks:
  - Generator G
  - Discriminator D
- Objective: Minimize the difference between real and generated data.

GAN Architecture

- The Generator G generates new data instances.
- The Discriminator D evaluates the authenticity of the generated instances.
- Both networks compete against each other, hence the term "adversarial."

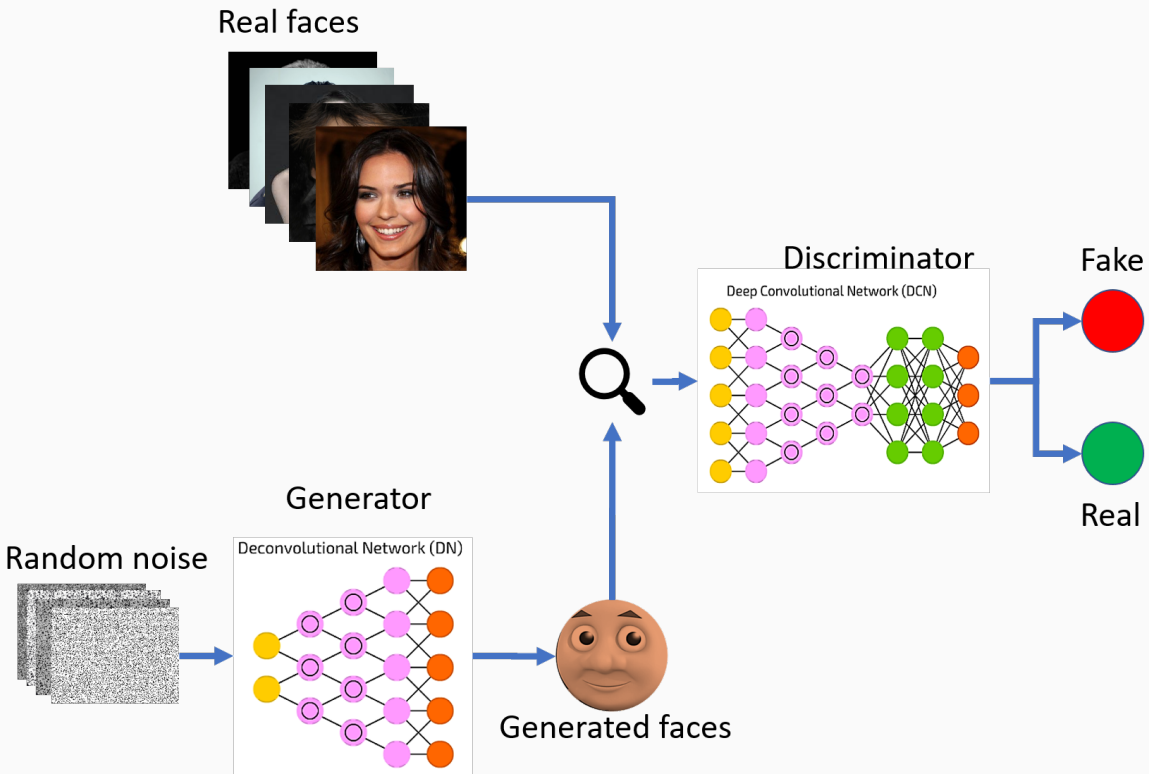

Mathematical Formulation

- Objective function:

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

- p_{data} is the distribution of real data and p_z is the prior distribution of input noise.
- The goal is to maximize the probability of correctly classifying real and generated samples.

Training Process

- GANs are trained using a two-step process:
  1. **Train the Discriminator D**:
     - Given real data, maximize D to correctly identify real vs generated data.
  2. **Train the Generator G**:
     - Given noise z, minimize $D(G(z))$ to improve the generation of realistic data.
- This process continues iteratively until convergence.

Loss Functions

- The loss function for the Discriminator:

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

- The loss function for the Generator:

$$L_G = -\mathbb{E}_{z \sim p_z}[\log D(G(z))]$$

- The generator aims to maximize the probability of the discriminator making a mistake.

Applications and Use Cases

- Image generation and enhancement
- Video generation
- Text-to-image synthesis
- Image super-resolution
- Semi-supervised learning
- Art generation and style transfer

Challenges in GANs

- Mode collapse: Generator produces a limited variety of samples.
- Training instability: Difficulty in balancing the generator and discriminator.
- Evaluation metrics: Lack of standardized metrics for evaluating GAN performance.
- High computational costs for training.

Future Directions

- Exploring new architectures (e.g., Progressive Growing GANs).
- Improved training methods (e.g., Wasserstein GANs).
- Better evaluation metrics for GANs.
- Applications in various domains, including healthcare, robotics, and gaming.

Conclusion

- GANs have revolutionized the field of generative modeling.
- They are powerful tools for generating realistic data and have numerous applications.
- Ongoing research is addressing challenges and exploring new possibilities.