Master 2 Computer Science

Advanced Neural Networks Architectures

Generative Models

Or how magic happens in AI

Part 2. Variational AutoEncoders

Akka Zemmari

Variational AutoEncoders (VAE)

Overview

- Introduction to Variational Autoencoders
- Variational Inference
- Variational Autoencoder Architecture
- Mathematical Formulation
- Applications and Use Cases

Variational AutoEncoders (VAE)

Whiteboard time!

Variational AutoEncoders (VAE)

Goal: Learn a generative model of data x.

- Probabilistic extension of AE: Encodes x into a distribution $z \sim q(z|x)$.
- Loss function combines reconstruction and KL Divergence.
- Objective: Maximize the Evidence Lower Bound (ELBO).
- Reparameterization trick for backpropagation.
- Applications: Image generation, semi-supervised learning, anomaly detection.
- Hands-on: Implement a VAE in PyTorch.
- Comparison with AE and GAN.

Variational Inference

- Variational inference is a technique to approximate complex probability distributions.
- It transforms the problem of inference into an optimization problem.
- Introduces a variational distribution $q(z|x)$ to approximate the true posterior $p(z|x)$.
- Objective: Minimize the Kullback-Leibler (KL) divergence:

$$D_{KL}(q(z|x)\|p(z|x)) = \int q(z|x) \log \frac{q(z|x)}{p(z|x)} dz$$

Variational Autoencoder Architecture

- An encoder that outputs the parameters of the variational distribution.
- A latent space z sampled from the variational distribution.
- A decoder that reconstructs the input from the latent space.

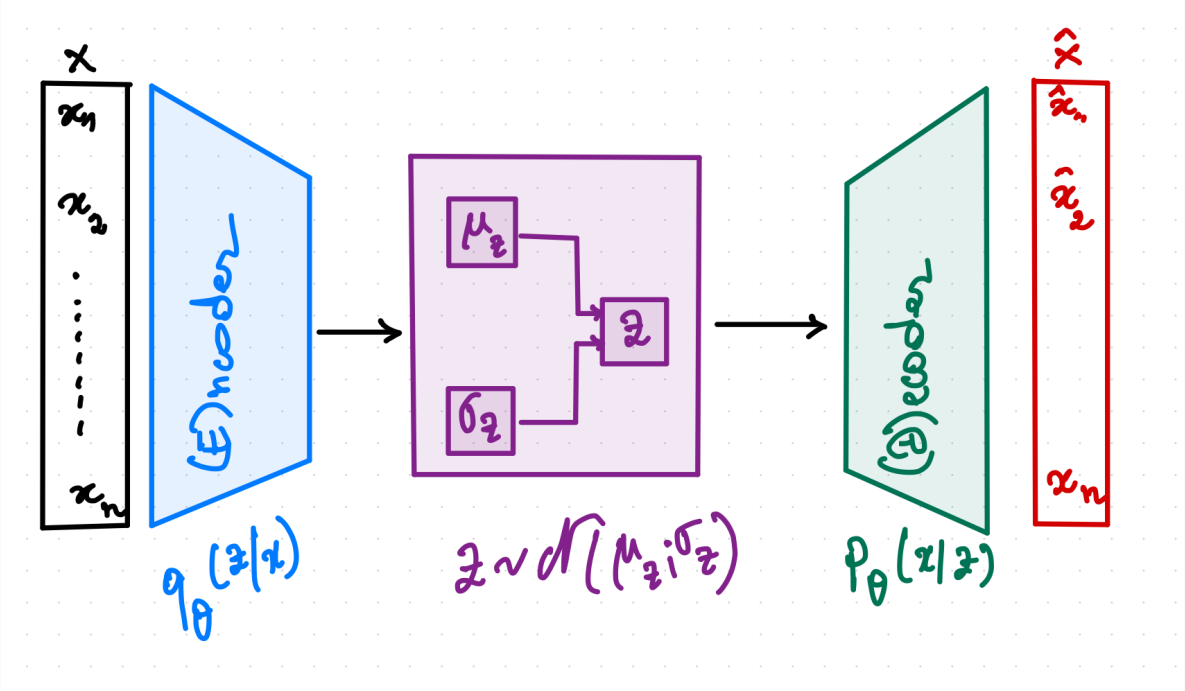

Figure 1: Variational Autoencoder Architecture

Mathematical Formulation

- Evidence Lower Bound (ELBO):

$$\mathcal{L}(\theta, \phi; x) = \mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x|z)] - D_{KL}(q_\phi(z|x) \| p(z))$$

- θ represents the parameters of the decoder and ϕ the encoder.
- Maximize ELBO to approximate the log-likelihood of data.
- Loss function:

$$\mathcal{L}_{\text{loss}} = -\mathcal{L}(\theta, \phi; x)$$

Reparameterization Trick

- To allow backpropagation through the stochastic sampling, we use the reparameterization trick:

$$z = \mu + \sigma \cdot \epsilon, \quad \epsilon \sim \mathcal{N}(0, I)$$

- Here, μ and σ are outputs of the encoder, and ϵ is noise.

Applications and Use Cases

- Image generation
- Semi-supervised learning
- Anomaly detection
- Data compression

Conclusion

- Variational Autoencoders combine the strengths of autoencoders and probabilistic graphical models.
- They enable powerful generative modeling and efficient inference.
- A versatile tool in machine learning and deep learning.