

Deep Learning

7.4 Variational Autoencoders (VAE)

Dr. Konda Reddy Mopuri
kmopuri@iittp.ac.in
Dept. of CSE, IIT Tirupati

Autoencoders

- ① Designed to reproduce input, especially reproduce the input from a learned encoding

Autoencoders

- ① Designed to reproduce input, especially reproduce the input from a learned encoding

Autoencoders


- ① Designed to reproduce input, especially reproduce the input from a learned encoding
- ② We attempted to project the data into the latent space and model it via a probability distribution

Autoencoders

- ① Designed to reproduce input, especially reproduce the input from a learned encoding
- ② We attempted to project the data into the latent space and model it via a probability distribution
- ③ This wasn't satisfying


Variational Autoencoders

- 1 Key idea is to make both Encoder and Decoder stochastic




Variational Autoencoders

- ① Key idea is to make both Encoder and Decoder stochastic
- ② Latent variable z is drawn from a probability distribution for the given input x



Variational Autoencoders

- ① Key idea is to make both Encoder and Decoder stochastic
- ② Latent variable z is drawn from a probability distribution for the given input x
- ③ Also, the reconstruction is chosen probabilistically from the sampled z



VAE Encoder

- ① Takes input and returns the parameters of a probability density (e.g. Gaussian, mean and covariance matrix)

VAE Encoder

- ① Takes input and returns the parameters of a probability density (e.g. Gaussian, mean and covariance matrix)
- ② We can sample this to get random values of the latent variable z

VAE Encoder

- ① Takes input and returns the parameters of a probability density (e.g. Gaussian, mean and covariance matrix)
- ② We can sample this to get random values of the latent variable z
- ③ NN implementation of the encoder gives (for every input x) a vector mean and a diagonal covariance

VAE Decoder

- ① Decoder takes the latent vector z and returns the parameters for a distribution

VAE Decoder

- ① Decoder takes the latent vector z and returns the parameters for a distribution
- ② $p_{\phi}(x/z)$ gives mean and variance for each pixel in the output

VAE Decoder

- ① Decoder takes the latent vector z and returns the parameters for a distribution
- ② $p_{\phi}(x/z)$ gives mean and variance for each pixel in the output
- ③ Reconstruction of x is via sampling

VAE loss function

- ① Loss for AE: l_2 distance between the input and its reconstruction

VAE loss function

- ① Loss for AE: l_2 distance between the input and its reconstruction
- ② In case of VAE: we need to learn parameters of two probability distributions

VAE loss function

- ① Loss for AE: l_2 distance between the input and its reconstruction
- ② In case of VAE: we need to learn parameters of two probability distributions
- ③ For each input x_i we maximize expected value of returning x_i (or, minimize the NLL)

$$-\mathbb{E}_{z \sim q_\theta(z/x_i)}[\log p_\phi(x_i/z)]$$

VAE loss function

$$-\mathbb{E}_{z \sim q_{\theta}(z/x_i)}[\log p_{\phi}(x_i/z)]$$

- ① Problem: Input images may be memorized in the latent space \rightarrow similar inputs may get different representations in z space

VAE loss function

$$-\mathbb{E}_{z \sim q_{\theta}(z/x_i)}[\log p_{\phi}(x_i/z)]$$

- ① Problem: Input images may be memorized in the latent space \rightarrow similar inputs may get different representations in z space
- ② We prefer continuous latent representations to give meaningful parameterization (e.g. smooth transition between digits)

VAE loss function

$$-\mathbb{E}_{z \sim q_{\theta}(z/x_i)}[\log p_{\phi}(x_i/z)]$$

- ① Problem: Input images may be memorized in the latent space \rightarrow similar inputs may get different representations in z space
- ② We prefer continuous latent representations to give meaningful parameterization (e.g. smooth transition between digits)
- ③ Solution: Force $q_{\theta}(z/x_i)$ to be close to a standard distribution (e.g. Gaussian)

VAE loss function


$$l_i(\theta, \phi) = -\mathbb{E}_{z \sim q_\theta(z/x_i)}[\log p_\phi(x_i/z)] + \mathbb{KL}(q_\theta(z/x_i) || p(z))$$

- ① First term promotes recovery, second term keeps encoding continuous (beats memorization)

VAE loss function

$$l_i(\theta, \phi) = -\mathbb{E}_{z \sim q_\theta(z/x_i)}[\log p_\phi(x_i/z)] + \text{KL}(q_\theta(z/x_i) || p(z))$$

① Problem: Differentiating over θ and ϕ



VAE loss function

$$l_i(\theta, \phi) = -\mathbb{E}_{z \sim q_\theta(z/x_i)}[\log p_\phi(x_i/z)] + \text{KL}(q_\theta(z/x_i) || p(z))$$

- ① Reparameterization: Draw samples from $N(0,1)$ \rightarrow doesn't depend on parameters

