Recitation: Homework 6
VAEs

10-418/10-618: ML for Structured Data

November 18, 2022

1. A simple example of a VAE

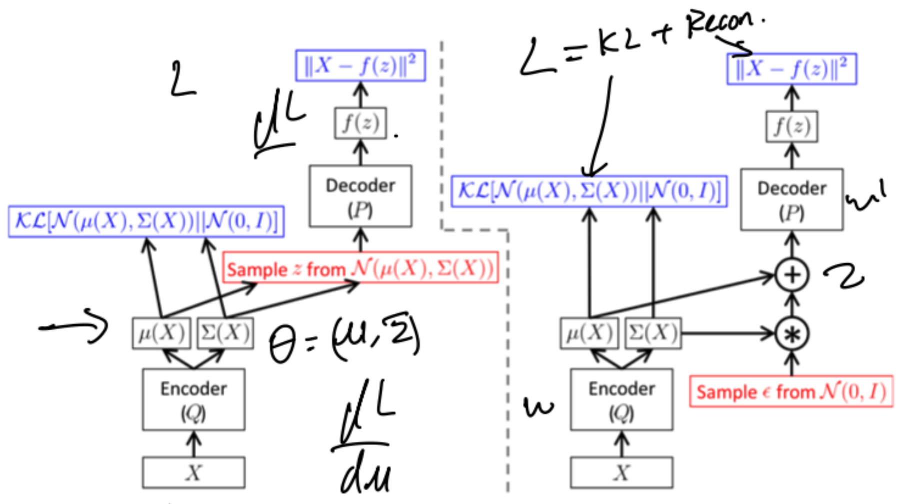

Figure 1: VAE structure; Right: reparameterization trick.

Figure 1 shows the structure of a VAE whose latent distribution is a Gaussian distribution parameterized by μ and Σ. Concretely, the input X will go through the encoder Q to produce the parameters of the Gaussian distribution. Then, the latent vector z sampled from the Gaussian distribution will be passed through the decoder P to reconstruct the input X. On the right side of Figure 1, we show the reparameterization trick: instead of sampling z from $\mathcal{N}(\mu, \Sigma)$, we sample $\epsilon \sim \mathcal{N}(0, I)$ and obtain $z = \mu + \Sigma^T \epsilon$.

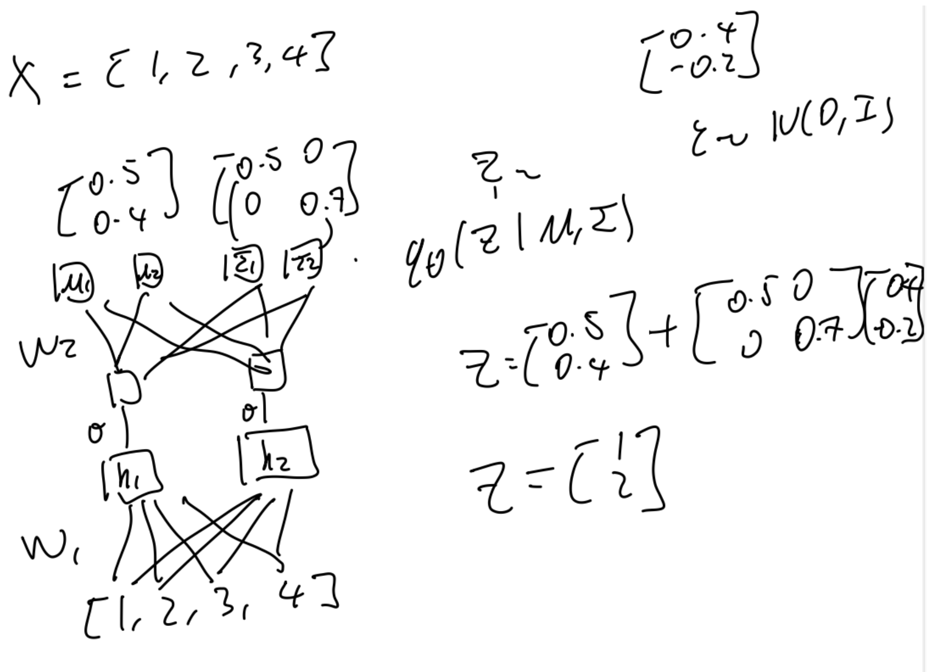

Figure 2: An example of VAE computation

Figure 2 gives an example of how to compute the parameters μ, Σ from a input of dimension 4 with a two-layer neural network as the encoder. Then, it shows how to use the reparameterization trick to obtain z from the sampled ϵ. Afterwards, z will be passed through a decoder with a similar architecture to output $f(z)$, which will be supervised to be close to X with the **reconstruction loss**.

$$\ell(\theta, \phi)$$
$$= -ELBO(\theta, \phi)$$
$$= \underbrace{-E_{z \sim q_\theta(z|x)}[\log p_\phi(x|z)]}_{\text{reconstruction error}} + \underbrace{KL(q_\theta(z|x) \,||\, p(z))}_{\text{regularization}}$$
$$\propto MSE[\hat{x}, x] - \frac{1}{2}\sum_{i=1}^{d}(1 + \log(\Sigma_i^2) - \mu_i^2 - \Sigma_i^2)$$

Figure 3: Decomposing ELBO

Last, we showed that the ELBO (Evidence lower bound) can be decomposed into two terms: a reconstruction term and a KL term, shown in Figure 3. In this way, we can supervise the encoder (parameterized by θ) and decoder (parameterized by ϕ) with these losses. Thanks to the reparameterization trick, we can back-propagate the gradient of the loss with respect to the weights of the networks (both the encoder and decoder) and update them via gradient descent.