Learning Latent Subspaces in Variational Autoencoders

Pierre-Aurelien Gilliot

Table of contents

Vanilla VAE

- x high dimensional
- Goal: sample from $p^*(\mathsf{x})$, get $p^*(z|\mathsf{x})$ (and z ?)
- Parameterization with NN:
  $p^*(\mathsf{x}|z) \approx p_\theta(\mathsf{x}|z) = \mathcal{D}(\mathsf{x}; \theta(z))$
  $p^*(z|\mathsf{x}) \approx q_\phi(z; \mathsf{x}) = \mathcal{N}(z; \mu_\phi(\mathsf{x}), \sigma_\phi(\mathsf{x}) I)$
- Optimizing the ELBO :

$$\begin{aligned}\log p_\theta(\mathsf{x}) &= \mathcal{L}_{\theta,\phi}(\mathsf{x}) + \mathcal{D}_{KL}(q_\phi(\mathsf{z} \mid \mathsf{x}) \mid\mid p_\theta(\mathsf{z} \mid \mathsf{x})) \\ &\geq E_{q_\phi(\mathsf{z}|\mathsf{x})}[\log p_\theta(\mathsf{x} \mid \mathsf{z})] - \mathcal{D}_{KL}(q_\phi(\mathsf{z} \mid \mathsf{x}) \mid\mid p_\theta(\mathsf{z}))\end{aligned}$$

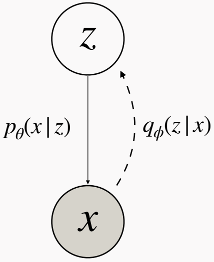

VAE-Demo

on Google colab, Tf 2.0

https://colab.research.google.com/drive/12_9dp3yJF6U_FM8CY98O9O5SIcDNMdV1?usp=sharing

- Interpretation of z?
- Dimension of z ?
- How to navigate z?

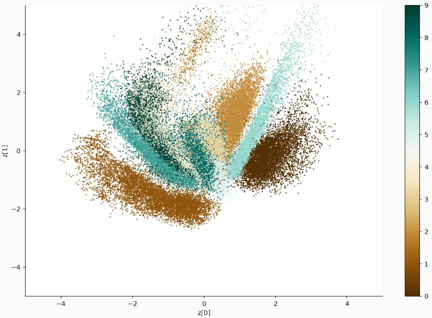

Conditional VAE

conditional VAE

- Conditioning on label y
- Objective:
  $E_{q_\phi(z|x,y)}[\log p_\theta(x \mid z, y)] - \mathcal{D}_{KL}(q_\phi(z \mid x, y) \mid\mid p_\theta(z, y))$
- Manipulate and control data
- Interpretable structure?

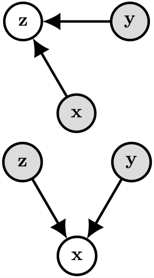

cVAE-Demo

on Google colab, Tf 2.0

https://colab.research.google.com/drive/1Cupg-5DS0GUikwJ0atEuMiZZPptiM0bI?usp=sharing