Generative Models: Implementations

32x32 CIFAR-10

Labeled Faces in the Wild

ECE324, Winter 2022

Michael Guerzhoy

For a single input case x:

$$L(x,\theta,\phi) = E_z[\log p_\theta(x|z)] - D_{KL}(q_\phi(z|x)||p_\theta(z))$$

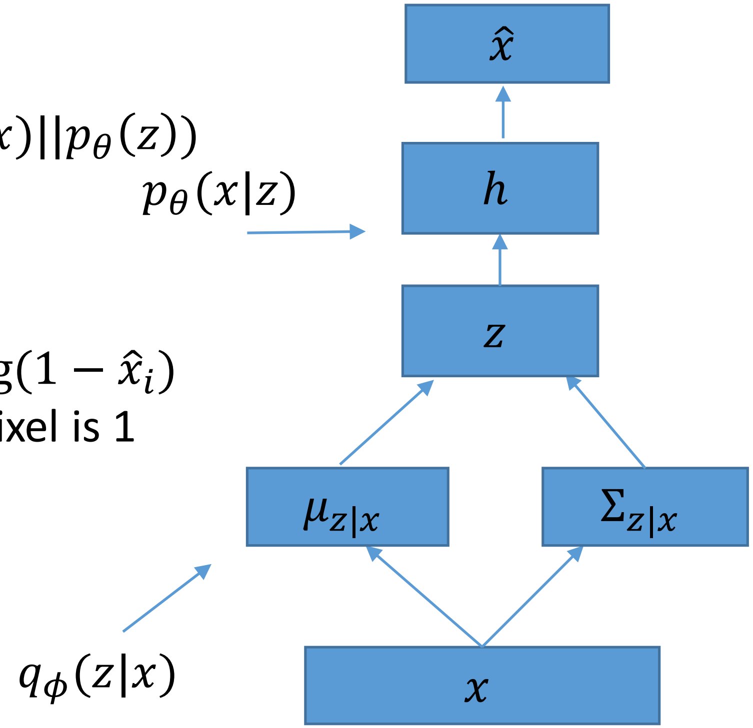

$$E_z[\log p_\theta(x_i|z)] \approx x_i \log \hat{x}_i + (1 - x_i) \log(1 - \hat{x}_i)$$

Pretend x_i is the probability that the i-th pixel is 1

$$p_\theta(z) = N(0, I), q_\phi(z|x) \sim N\left(\begin{matrix}\mu_1(x)\\ \dots\\ \mu_k(x)\end{matrix}, \begin{pmatrix}\sigma_1(x) & & \\ & \dots & \\ & & \sigma_k(x)\end{pmatrix}\right)$$

$$D_{KL}(q_\phi(z|x)||p_\theta(z)) = \frac{1}{2}\left[\sum_j \log \sigma_j - k + \sum_j \mu_j^2 + \sum_j \sigma_j\right]$$

Reparametrization trick

- Want to learn to sample good z's using
$$z \sim N\left(\mu_{z|x}, \sigma_{\sigma|x}\right)$$
- The z's sampled depend on μ and Σ, but we cannot differentiate with respect to them
- Trick:
$$\epsilon \sim N(0, 1)$$
$$z = \mu_{z|x} + \epsilon \sigma_{z|x}$$
- Now can differentiate wrt μ and Σ

Training GANs

for number of training iterations **do**

for k steps **do**

- Sample minibatch of m noise samples $\{\boldsymbol{z}^{(1)}, \ldots, \boldsymbol{z}^{(m)}\}$ from noise prior $p_g(\boldsymbol{z})$.
- Sample minibatch of m examples $\{\boldsymbol{x}^{(1)}, \ldots, \boldsymbol{x}^{(m)}\}$ from data generating distribution $p_{\text{data}}(\boldsymbol{x})$.
- Update the discriminator by ascending its stochastic gradient:

$$\nabla_{\theta_d} \frac{1}{m} \sum_{i=1}^{m} \left[\log D_{\theta_d}(x^{(i)}) + \log(1 - D_{\theta_d}(G_{\theta_g}(z^{(i)}))) \right]$$

end for

- Sample minibatch of m noise samples $\{\boldsymbol{z}^{(1)}, \ldots, \boldsymbol{z}^{(m)}\}$ from noise prior $p_g(\boldsymbol{z})$.
- Update the generator by ascending its stochastic gradient (improved objective):

$$\nabla_{\theta_g} \frac{1}{m} \sum_{i=1}^{m} \log(D_{\theta_d}(G_{\theta_g}(z^{(i)})))$$

end for