LECTURE 17

Generative Models

Recall classifier based on CNN

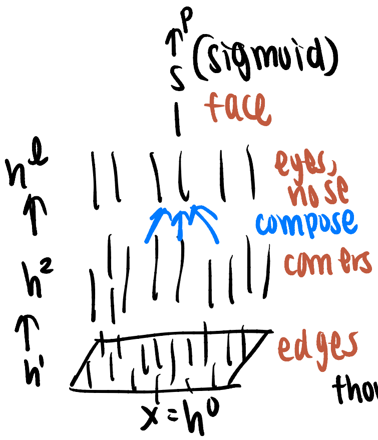

Discriminator - tell if image face or not

bottom-up structure

$$h^l_{ij} = \sum_{\Delta i, \Delta j} w^l_{\Delta i, \Delta j} h^{l-1}_{i+\Delta i, j+\Delta j} + b^l + \text{Rectify}$$

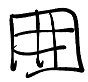

thought vector

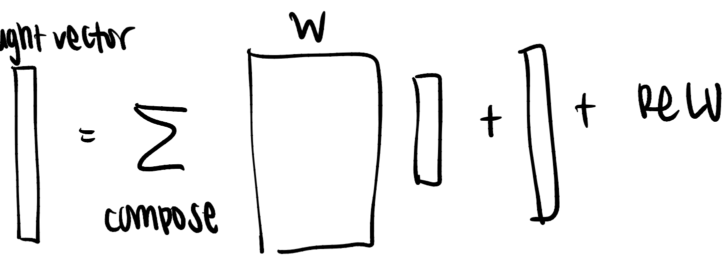

- thought vectors at each layer
- more complex
- see bigger patterns

Top down

Generator
↳ can generate images

$h^L = z \mid z \sim N(0, I)$

still a CNN

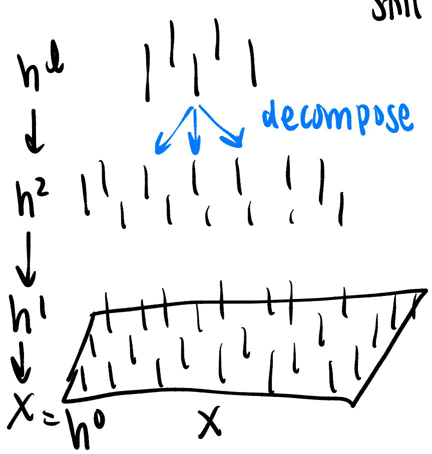

$$h^l_{ij} = \sum_{\Delta i, \Delta j} w^l_{\Delta i, \Delta j} h^{l+1}_{i+\Delta i, j+\Delta j} + b^l + \text{Rectify}$$

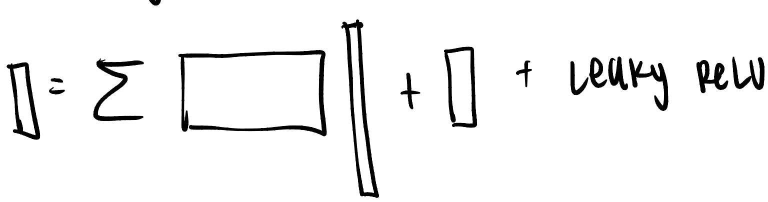

Generative adversarial networks (GAN)

real	1 ⋮ n	x_i	$y_i = 1$
fake	1 ⋮ n	$\tilde{x}_i$	$\tilde{y}_i = 0$

Discriminator

$D(x) = Pr(y = 1 \mid x)$

- train like classification
- logistic regression / max likelihood

Generator

$\tilde{x} = G(\tilde{z})$

$\tilde{z} \sim N(0, I)$

Maximum likelihood for D

$$\frac{1}{n}\sum_{i=1}^{n}\log P(y_i=1 \mid x_i)+\frac{1}{n}\sum_{i=1}^{n}\log P(\tilde{y}_i=0 \mid \tilde{x}_i)$$

$$=\frac{1}{n}\sum_{i=1}^{n}\log D(x_i)+\frac{1}{n}\sum_{i=1}^{n}\log(1-D(\tilde{x}_i))$$

want to assign high probs to real examples
want to assign low probs to fake examples

$$V=\frac{1}{n}\sum_{i=1}^{n}\log D(x_i)+\frac{1}{n}\sum_{i=1}^{n}\log(1-D(G(\tilde{z}_i)))$$

learn D by maximizing value (log likelihood)

learn G by minimizing value (log likelihood)

Adversarial Game (zero sum)

train 2 networks to compete w/ each other
⇒ create very powerful generator

$$\min_G \max_D V$$

reach Nash equilibrium so sol'n to min max = sol'n to max min
↳ saddle point

G (thief) increase $D(G(\tilde{z}_i))$ to fool D (police) → tell apart fake & real examples

increase prob. of fake examples as real

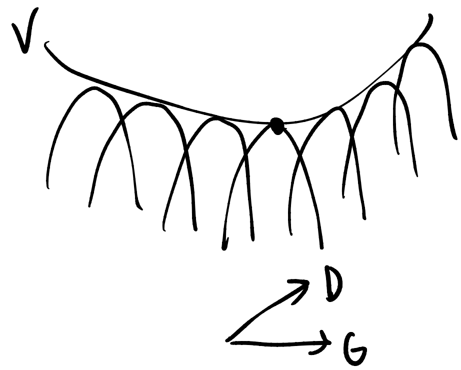

Wasserstein GAN

$$\min_G \max_f \left[\frac{1}{n}\sum_{i=1}^{n} f(x_i)-\underbrace{\frac{1}{n}\sum_{i=1}^{n} f(G(\tilde{z}_i))}_{\text{G try to max this}}\right]$$

$\tilde{x}_i$

G → actor
f → critic

actor: generate fake examples to get high score

critic: outputs a score; criticize fake examples, high score to real examples, low score to fake examples

regularized

Generator

$X = G(z)$

z: latent code / embedding

encoder $z = \mu_\phi(x)$

decoder $X = G_\theta(z)$

X

Auto-encoder

loss function

$$\min_{\theta,\phi} \frac{1}{n}\sum_{i=1}^{n} \left| x_i - G_\theta(\underbrace{\mu_\phi(x_i)}_{z_i}) \right|^2$$

z is much lower dimension than x

bottleneck

dimension reduction

manifold assumption

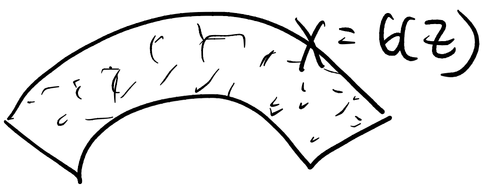

prob model

$z \sim N(0, I)$

$x = G(z) + \varepsilon \quad , \quad \varepsilon \sim N(0, \sigma^2 I)$

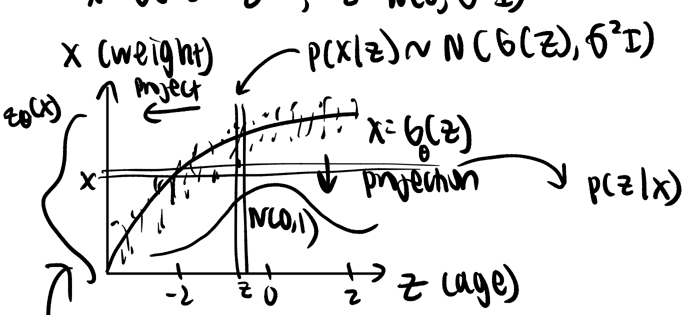

marginal distribution of x

$z = \mu_\phi(x) + \tau e, \quad e \sim N(0, I)$

$N(\mu_\phi(x), \tau^2)$

$q_\phi(z|x) \approx p(z|x)$

encoder

$z \sim p(z) \sim N(0, I)$

$p_\theta(x|z) \sim N(G_\theta(z), \sigma^2 I)$

decoder

X

$X = G_\theta(z) + \varepsilon, \quad \varepsilon \sim N(0, \sigma^2 I)$

variational Auto-encoder (VAE)

$$\max_{\theta,\phi} \frac{1}{n} \sum_{i=1}^{n} \left[\log p_\theta(x_i) - D_{KL}(q_\phi(z_i|x_i) \| p_\theta(z_i|x_i)) \right]$$

$D_{KL}(q \| p)$: kullback-leibler divergence