$D = \{(\bar{x}, y)\}$ ML model

Generative

$\boxed{p(\bar{x}, y)}$

Discriminative

$\boxed{p(y|\bar{x})}$

$f: \bar{x} \to y$

$p(y|\bar{x}) = \frac{p(\bar{x}, y)}{p(\bar{x})} \propto p(\bar{x}, y)$

Logistic regression

$p(y|\bar{x}) = \mathrm{softmax}(\bar{\theta}^T \bar{x})$

$p(y=k|\bar{x}) = \frac{e^{\bar{\theta}_k^T \bar{x}}}{\sum_\ell e^{\bar{\theta}_\ell^T \bar{x}}} \propto e^{\bar{\theta}_k^T \bar{x}}$

Naive Bayes

$p(\bar{x}, y) = p(y) \cdot \prod_{i=1}^{n} p(x_i | y)$

$\log p(\bar{x}, y) = \log p(y=k) + \sum \log p(x_i = m | y = k) =$

$\boxed{\theta_k, \; k = 1..K} \qquad \boxed{\theta_{imk}}$

$\log p(y=k|\bar{x}) =$

$= \mathrm{const} + \bar{\theta}_k^T \bar{x}$

$= \sum_{\ell=1}^{K} \theta_\ell [y=k] + \sum_i \sum_\ell \sum_s \theta_{is\ell} \cdot [x_i = s, y = \ell]$

features

$p(\bar{x}) \qquad p(D|\theta) \longrightarrow \max$

$\theta = \arg\max \prod_{\bar{x} \in D} p(\bar{x}|\theta) = \arg\max \boxed{\sum_{\bar{x} \in D} \log p(\bar{x}|\theta)}$

$KL(p\|q) \neq KL(q\|p)$

Pdata

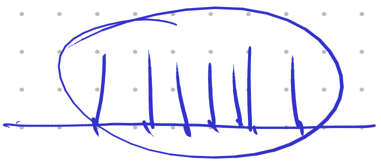

D

$\approx$

$p_{model}(x|\theta)$

Kullback-Leibler divergence

$KL(p\|q) = \int p(x) \log \frac{p(x)}{q(x)} dx$

$\geq 0 \quad , \quad = 0 \iff p = q$

$\bar{\theta} \xleftarrow{} \min$

$KL(p_{data} \| p_{model}) = \int p_{data}(\bar{x}) \cdot \log \frac{p_{data}(\bar{x})}{p_{model}(\bar{x})} d\bar{x} =$

$= \int p_{data}(\bar{x}) \log p_{data}(\bar{x}) dx - \int p_{data}(\bar{x}) \boxed{\log p_{model}(\bar{x})} d\bar{x}$

$= \mathrm{const} - \frac{1}{N} \sum_{\bar{x} \in D} \log p_{model}(\bar{x}) \quad \bar{\theta} \longrightarrow \max$

① Explicit density models

$p(x)$

1.1. Tractable density

$$p(\bar{x}) = p(x_1, x_2, \dots, x_n) = p(x_1)\,p(x_2|x_1)\,p(x_3|x_1, x_2) \dots p(x_n|x_1 \dots x_{n-1})$$

$$p(x_i|x_1 \dots x_{i-1})$$

NB: $p(\bar{x}, y) = p(y) \cdot p(x_1|y) \dots p(x_n|y)$

RNN:

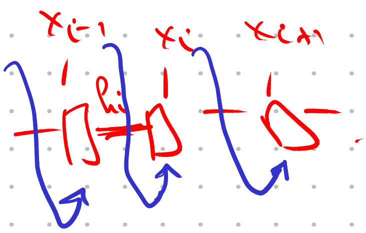

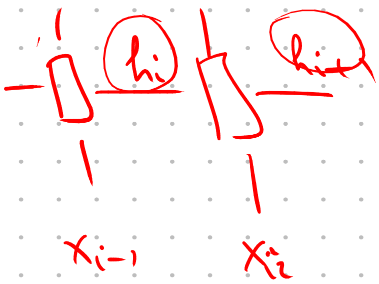

$p(x_i | x_1 \dots x_{i-1}) =$

$p(x_i | h_i, x_{i-1})$

$p(h_i | h_{i-1}, x_{i-2})$

$p(h_{i-1} | h_{i-1}$

1.2. Approximate density

$p(\bar{x}) \approx q(\bar{x})$

VAE

Variational autoencoder

② Implicit density models / GAN / $\cancel{p(x)}$

$\sim Unif([0,1])$

$\cancel{D = \{(\bar{z}, \bar{x})\}}$

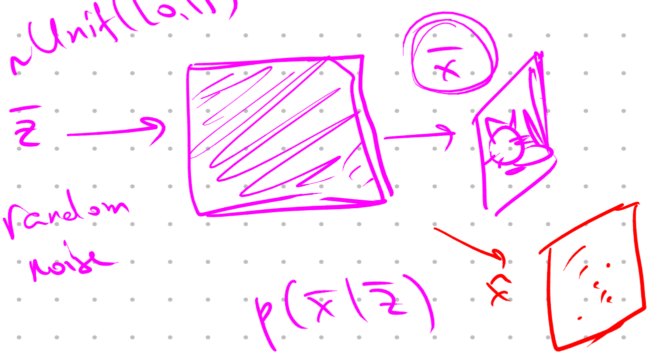

$p(\bar{x}|\bar{z})$

$L(\hat{x})$ = „catness"

Generative Adversarial Networks

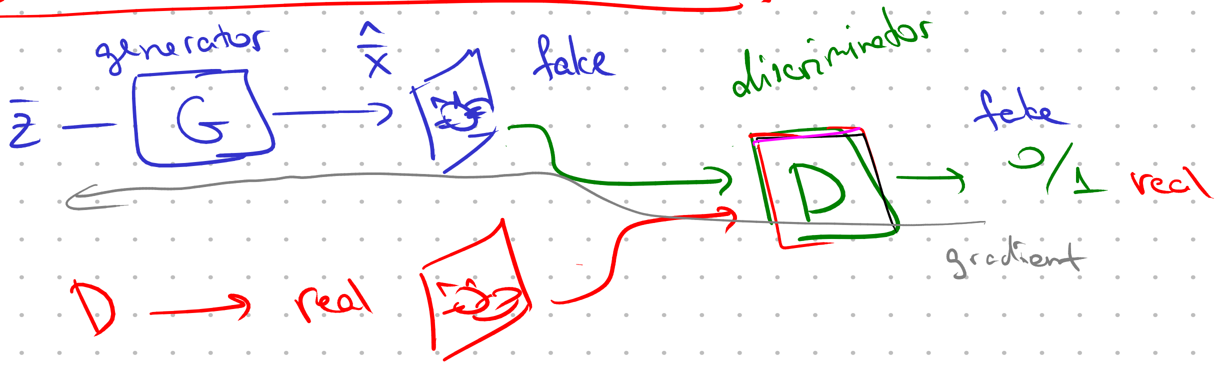

$$L_D(\bar{x}) = \mathbb{E}_{\bar{x}\in Real}[\log D(\bar{x})] + \mathbb{E}_{\bar{x}\in fake}[\log(1-D(\bar{x}))] \xrightarrow[\theta_D]{} max$$

$$L_D = \underbrace{\mathbb{E}_{Real}[\log D(\bar{x})] + \mathbb{E}_{\bar{z}}[\log(1-D(G(\bar{z})))]}_{V(D,G)} \xrightarrow[\theta_D]{} max$$

$$L_G = -L_D \xrightarrow[\theta_G]{} min$$

$$\min_G \underbrace{\max_D V(D,G)}$$

$$D^*(\bar{x}) = \frac{p_{data}(\bar{x})}{p_{data}(\bar{x}) + p_G(\bar{x})}$$

$$G^*(\bar{x}) = p_{data}(\bar{x})$$

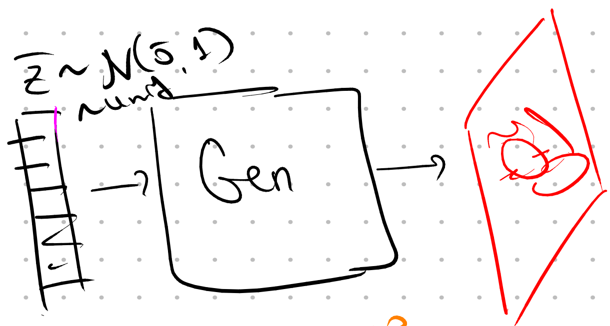

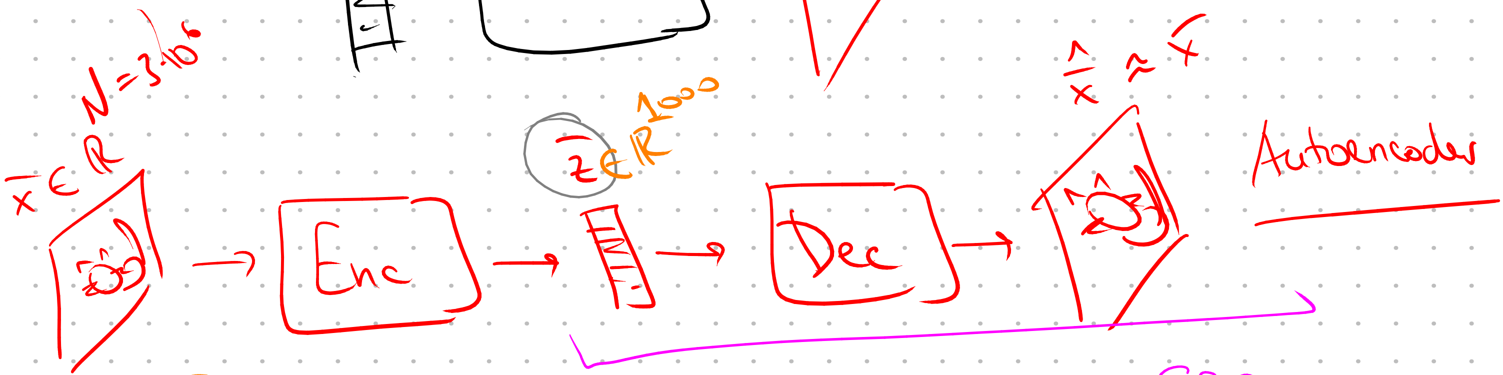

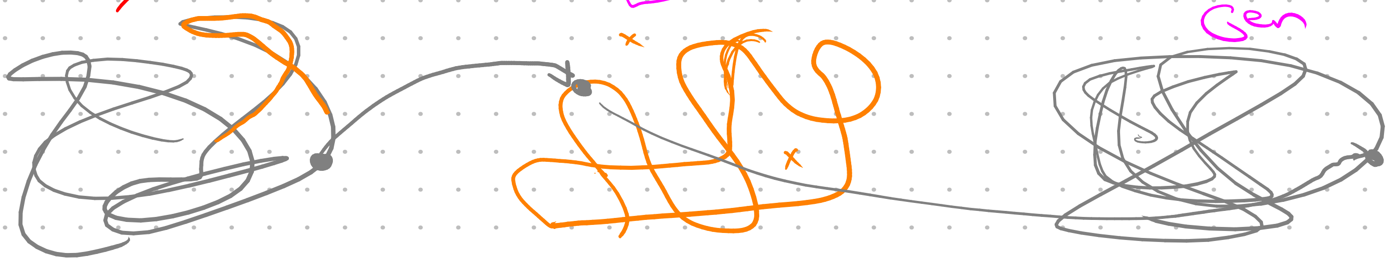

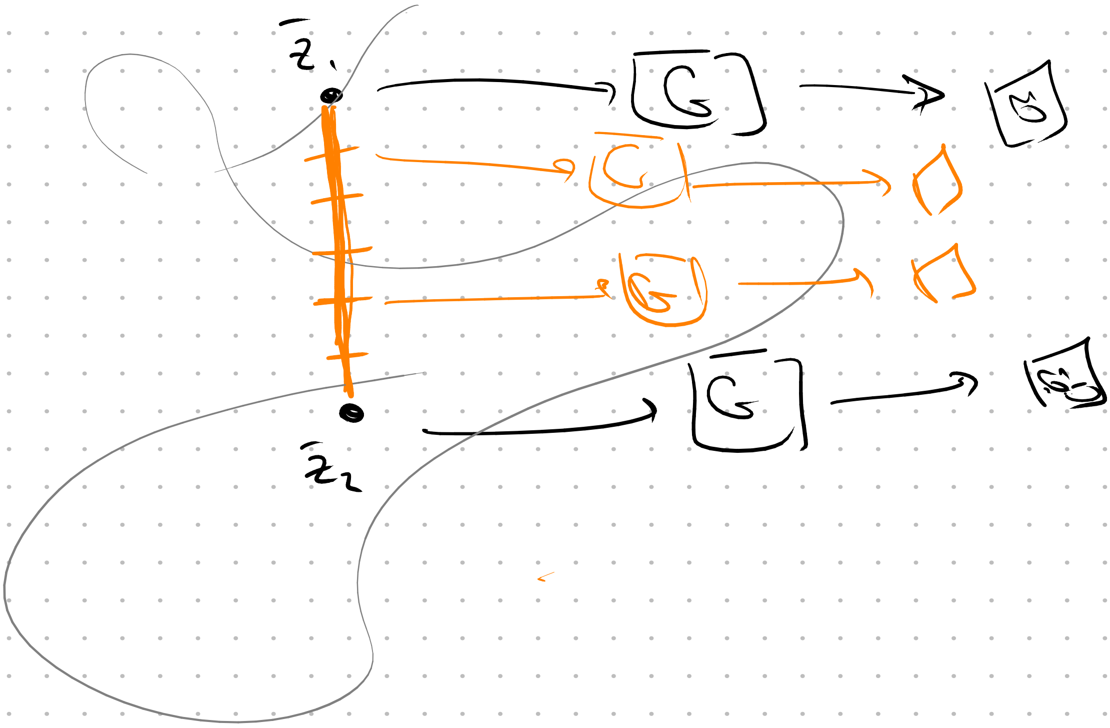

Adversarial autoencoder

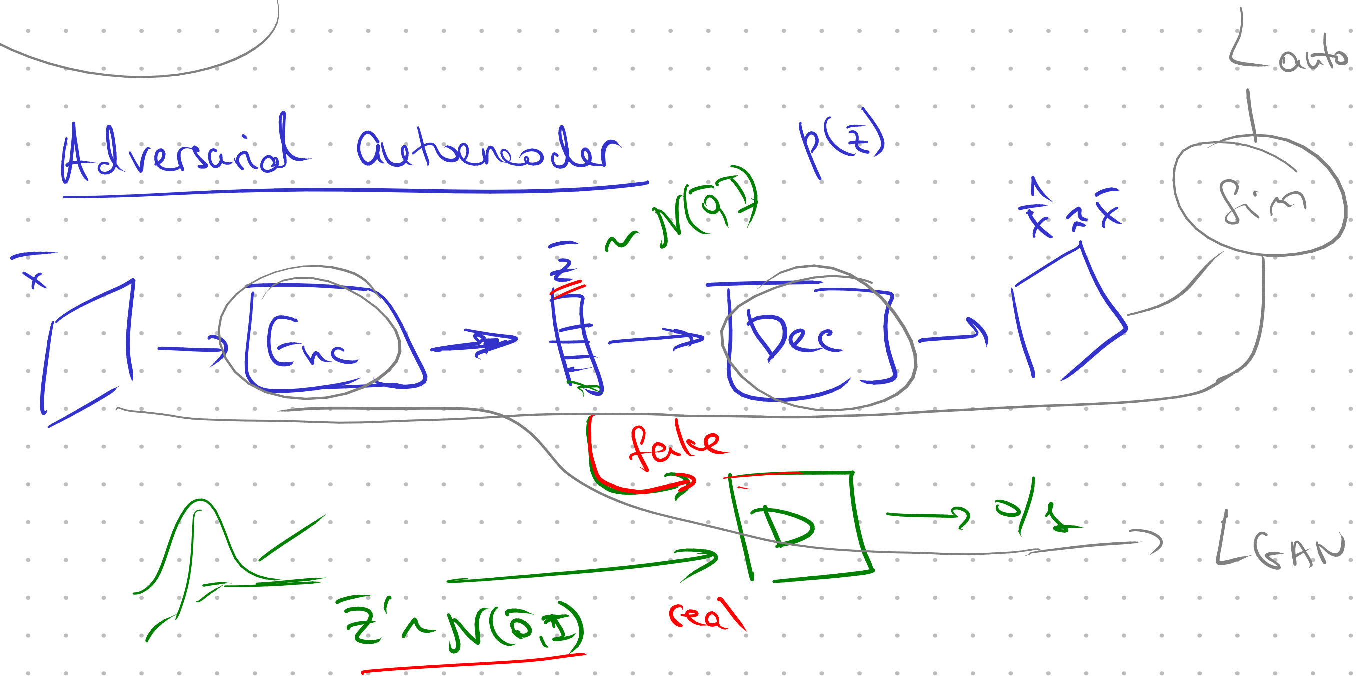

$$L_{AE} = L_{auto} + \lambda \cdot L_{GAN}$$