$\mathcal{D}$

Machine Learning → Reinforcement Learning

Agent $\xrightarrow{a_t}$ Env, Env $\xrightarrow{r_t,\, s_{t+1}}$ Agent

Supervised learning

$D = \{(\bar{x}_n, y_n)\}_{n=1}^{N}$

Regression: $y_n \in \mathbb{R}$

Classification: $y_n \in \{C_1, \ldots, C_K\}$, $p(C_1) \ldots p(C_K)$

Learning to rank

Unsupervised learning

$D = \{\bar{x}_n\}_{n=1}^{N}$

Clustering

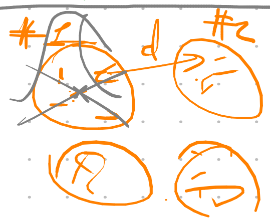

Dimensionality reduction

Anomaly detection

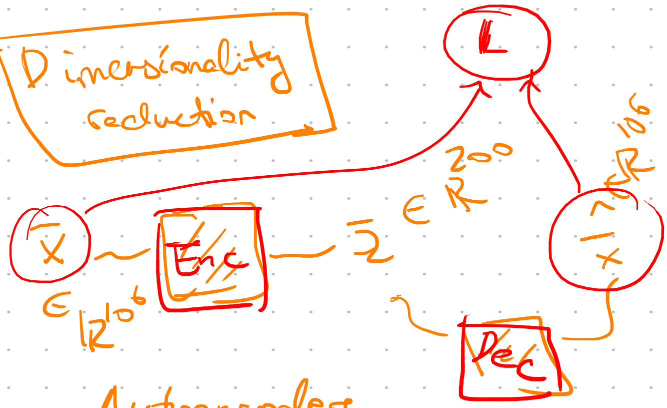

Autoencoders

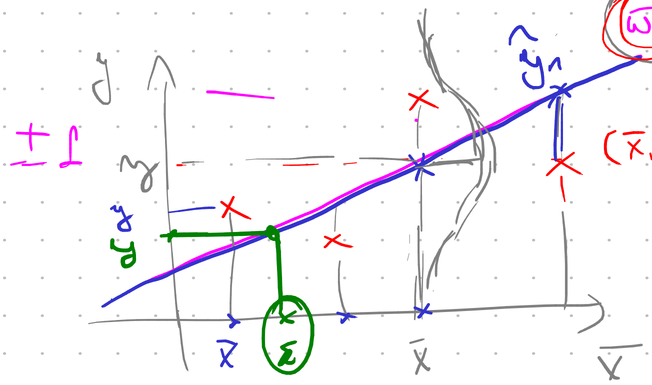

$\bar{w}^T \bar{x} = \hat{y}$

$\hat{y}_n \approx y_n$

$= \bar{w}^T \bar{x}_n$

MHK MSE

$$\sum_n (y_n - \hat{y}_n)^2 \longrightarrow \min$$

$(\quad)^4$

e^1

$\mathcal{D}$, $\bar{\theta}$

posterior, prior, likelihood

$$p(\bar{\theta} \mid D) = \frac{p(\bar{\theta})\, p(D \mid \bar{\theta})}{p(D)}$$

$$p(D) = \int p(\bar{\theta})\, p(D \mid \bar{\theta})\, d\bar{\theta}$$

1) $p(D \mid \bar{\theta}) \xrightarrow{\bar{\theta}} \max$ $\bar{\theta}_{ML}$

2) $p(\bar{\theta} \mid D) \xrightarrow{\bar{\theta}} \max \iff p(\bar{\theta})\, p(D \mid \bar{\theta}) \xrightarrow{\bar{\theta}} \max$ $\bar{\theta}_{MAP}$

$$p(D \mid \bar{\theta}) = \prod_{n=1}^{N} p(d_n \mid \bar{\theta})$$

$D = \{d_n\}_{n=1}^{N}$

$$p(y \mid x, \bar{\theta}) = \prod_{n=1}^{N} p(y_n \mid \bar{x}_n, \bar{\theta})$$

$$p(D|\bar w) = \prod_{n=1}^{N} \mathcal{N}(y_n \mid \bar w^T \bar x_n, \sigma^2) = \prod_{n=1}^{N} \frac{1}{\sqrt{2\pi\sigma^2}}\, e^{-\frac{1}{2\sigma^2}(y_n - \bar w^T \bar x_n)^2} \xrightarrow[\bar w]{} \max$$

$$\log p(D|\bar w) = \underbrace{-\frac{N}{2}\log(2\pi\sigma^2)} - \frac{1}{2\sigma^2}\sum_{n=1}^{N} \frac{1}{\sigma_n^2}(y_n - \bar w^T \bar x_n)^2 \xrightarrow[\bar w]{} \max$$

$$\xrightarrow[\bar w]{} \min$$

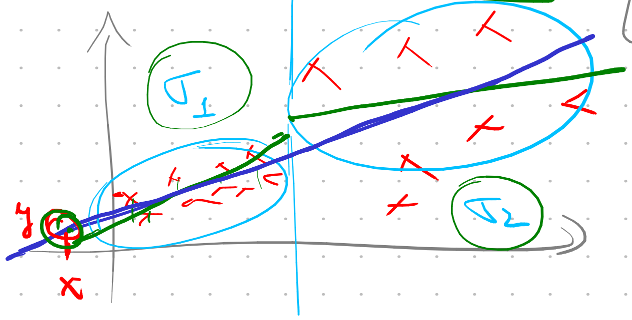

$$\frac{\partial(\ldots)}{\partial\theta} = n\theta^{n-1}(1-\theta)^m - m\theta^n(1-\theta)^{m-1} = 0$$

$\theta = p(\text{орёл})$ $p(\text{решка}) = 1 - \theta$

$$\theta^{n-1}(1-\theta)^{m-1}\big(n(1-\theta) - m\theta\big) = 0$$

$$p(\underbrace{hhtht}_{n \times h,\ m \times t}|\theta) = \theta^n(1-\theta)^m \xrightarrow[\theta]{} \max$$

$$\boxed{\theta_{ML} = \frac{n}{n+m}}$$

$$p(h|\theta) = \theta \to \max$$

$$\theta_{ML} = 1$$

prior

$$p(\theta|\text{heads}) = \frac{p(\theta)\,\theta^{550}(1-\theta)^{450}}{p(\text{heads})}$$

$$p(\theta) \propto \theta^{\alpha-1}(1-\theta)^{\beta-1}$$

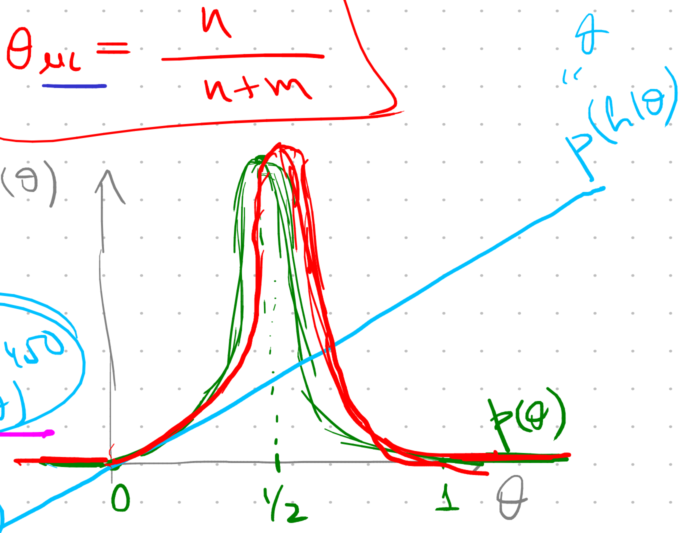

$$p(\bar w) = \mathcal{N}\left(\bar w \mid \bar 0, \frac{1}{\sigma_0^2} \cdot I\right) \qquad \begin{pmatrix} \sigma_0^2 & 0 \\ 0 & 1/\sigma_0^2 \end{pmatrix}$$

$$w_i \sim \mathcal{N}(w_i | 0, \sigma_0^2) \qquad p(w_i) \propto e^{-\text{Const}\cdot|w_i|}$$

$$\boxed{p(\bar\theta|D) \xrightarrow{\theta} \max}$$

$$\log p(\bar w|D) = \text{Const} + \log p(D|\bar w) + \log p(\bar w) = \text{Const} - \frac{1}{2\sigma^2}\sum_n (y_n - \bar w^T \bar x_n)^2 + \sum_{i=1}^{d}\left(\cancel{-\frac{1}{2}\log(2\pi\sigma_0^2)} - \frac{1}{2\sigma_0^2} \cdot w_i^2\right)$$

$$\log p(\bar w|D) = \text{const} - \frac{1}{2}\left(\frac{1}{\sigma^2}\sum_n \left(y_n - \bar w^T \bar x_n\right)^2 + \frac{1}{\sigma_0^2}\sum_i w_i^2\right) \xrightarrow[\bar w]{} \max$$

$\bar w \to \min$

$\sum_i |w_i|$

3) Predictive distribution

$$p(\bar x|D) = \int p(\bar x, \bar\theta|D)\,d\bar\theta = \int p(\bar x|\bar\theta, \cancel{D})\, p(\bar\theta|D)\,d\bar\theta =$$

$$= \int \underbrace{p(\bar x|\bar\theta)}_{\text{likelihood}}\, \underbrace{p(\bar\theta|D)}_{\text{posterior}}\,d\bar\theta = E_{p(\bar\theta|D)}\left[p(\bar x|\bar\theta)\right]$$

?

$p(y|\bar x, \bar\theta)$

θ_{MAP}

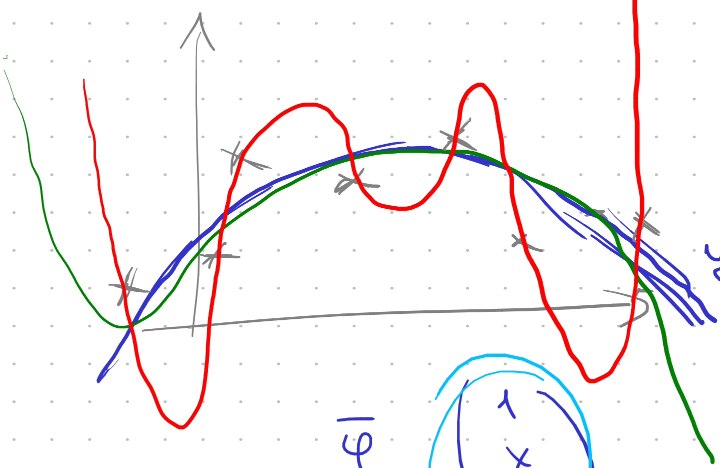

$y = \bar w^T \bar x$

$y = w \cdot x$

$$y = w_0 + w_1 x + w_2 x^2 = \bar w^T \begin{pmatrix} 1 \\ x \\ x^2 \end{pmatrix}$$

$$x \xrightarrow{\bar\varphi} \begin{pmatrix} 1 \\ x \\ x^2 \end{pmatrix}$$

SIFT

$\bar x$ — Feature extraction — $\bar\varphi(\bar x)$ — linear — y

softmax — $p(C_1) \dots p(C_k)$

$\bar x \in \mathbb{R}^{100} \rightsquigarrow y$

$\cancel{w^T x} \rightsquigarrow y$

$$\bar x = \begin{pmatrix} x_1 \\ \dots \\ x_{100} \end{pmatrix} \xrightarrow{\bar\varphi} \begin{pmatrix} x_1^2 \\ x_1 x_2 \\ \dots \\ x_{100}^2 \\ x_1 \\ \dots \\ x_{100} \end{pmatrix} + d$$

$$\frac{d(d-1)}{2} + d = \boxed{\frac{d(d+1)}{2}}$$

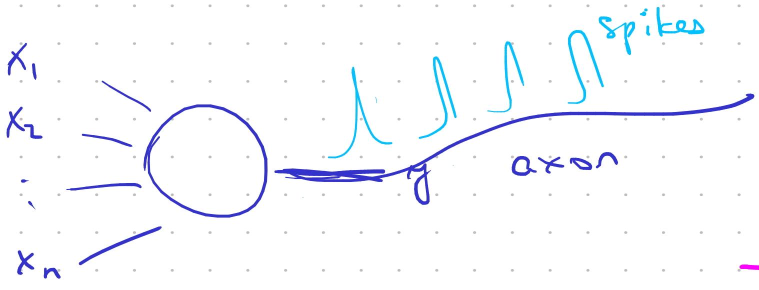

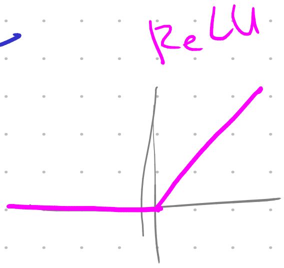

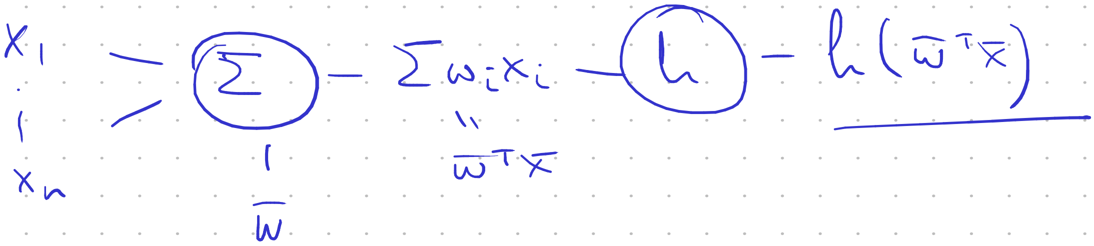

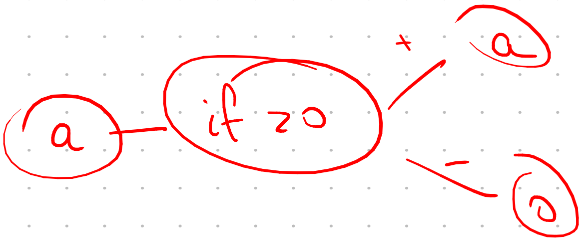

$$\text{ReLU}(a) = \max(a, 0)$$

$$\text{Swish}(a) = a \cdot \sigma(\beta \cdot a) =$$

$$= \frac{a}{1 + e^{-\beta a}}$$