ML → Reinforcement

Supervised | Unsupervised

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

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

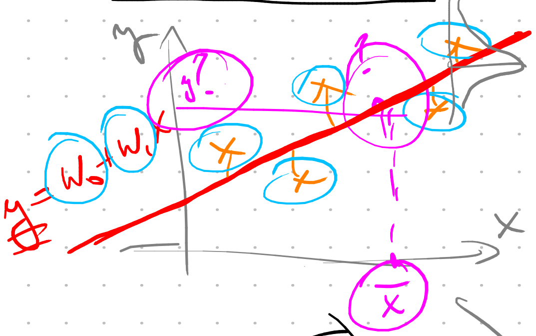

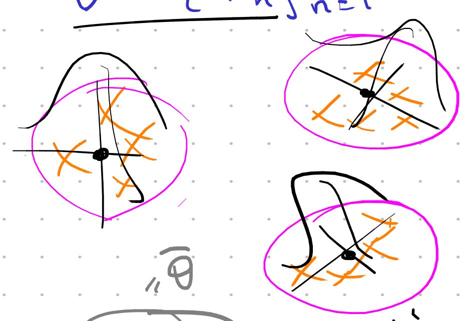

$$p(y \mid \bar{x}, \bar{\theta}) = \mathcal{N}(y \mid \bar{x}^T \bar{w}, \sigma^2)$$

$$p(\bar{x} \mid \bar{\alpha}, \bar{\mu}, \bar{\Sigma}) = \sum_{k=1}^{K} \alpha_k \mathcal{N}(\bar{x} \mid \bar{\mu}_k, \Sigma_k)$$

$$y = f_{\bar{\theta}}(\bar{x}) + \varepsilon_{\bar{\theta}}$$

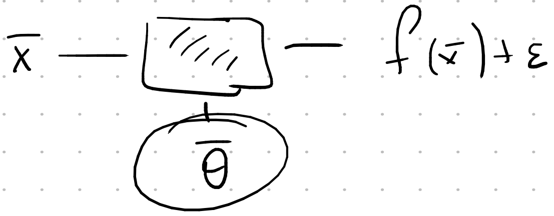

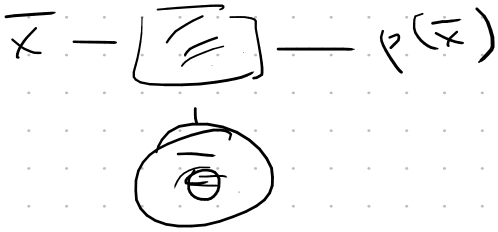

$$p(x, y) = p(x)\,p(y \mid x) = p(y)\,p(x \mid y)$$

$$p(y \mid x) = \frac{p(y)\,p(x \mid y)}{p(x)} \quad \text{— Bayes Rule}$$

posterior = prior × likelihood

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

$p(D)$ evidence prob

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

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

1) $\bar{\theta}_{ML} = \underset{\bar{\theta}}{\arg\max}\ p(D \mid \bar{\theta})$ — max likelihood

2) $\bar{\theta}_{MAP} = \underset{\bar{\theta}}{\arg\max}\ p(\bar{\theta} \mid D) = \underset{\bar{\theta}}{\arg\max}\ p(\bar{\theta})\,p(D \mid \bar{\theta})$ — maximum a posteriori

3) $$\underbrace{p(\bar{x} \mid D)}_{\text{predictive distrib.}} = \int p(\bar{x}, \bar{\theta} \mid D)\,d\bar{\theta} = \int \underbrace{p(\bar{x} \mid \bar{\theta}, \cancel{D})}_{\text{likelihood}}\ \underbrace{p(\bar{\theta} \mid D)}_{\text{posterior}}\,d\bar{\theta} = \mathbb{E}_{p(\bar{\theta} \mid D)}\left[p(\bar{x} \mid \bar{\theta})\right]$$