11/15/22

Baye's Rule

$$\begin{tikzcd} X \arrow[d, "①"] \\ Y \arrow[u, bend left, dashed, "②"] \end{tikzcd}$$

Solving equation

$$f(x) = y$$

$$\begin{tikzcd} X \sim p(x) \arrow[d, "\text{forward}: p(y|x)"] \\ Y \arrow[u, bend left, dashed, "\text{inverse}: p(x|y)"] \end{tikzcd}$$

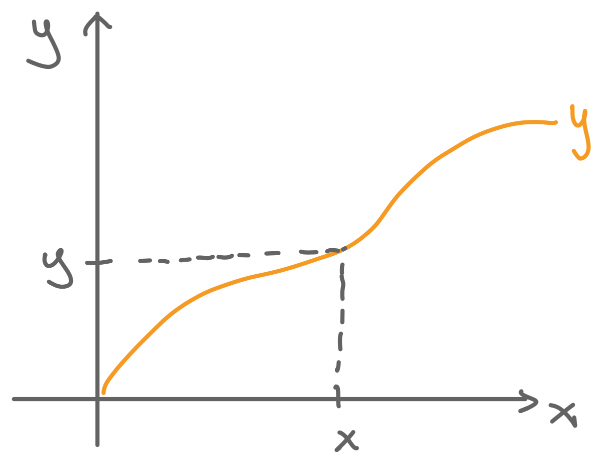

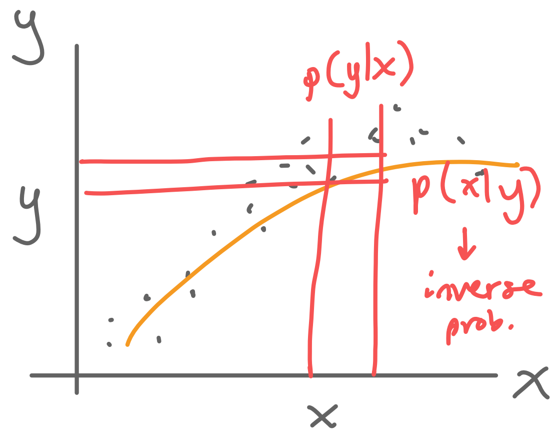

Factor Analysis

$$Z \to X$$

$$X = WZ + \varepsilon$$

embedding 10x3 matrix into vector

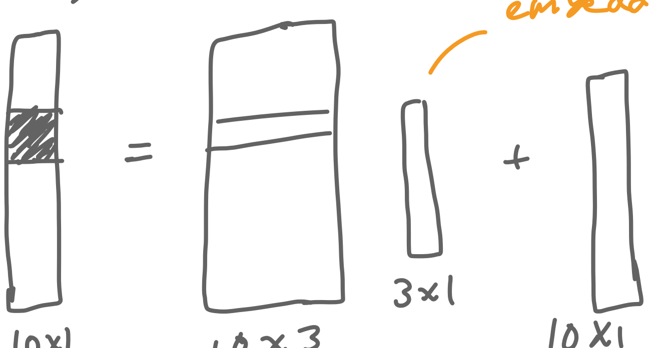

$\varepsilon \perp Z$ $\quad Z \sim N(0, I_3)$ $\quad \varepsilon \sim N(0, \sigma^2 I_{10})$

$$[z \mid x] \sim N(\qquad)$$

parameters $\theta = (w, \sigma^2)$

$$P_\theta(x) = \int P(z) P_\theta(x \mid z)\, dz$$

density estimation

$$\underset{10\times 1}{\boxed{X_1}} \quad \underset{10\times 1}{\boxed{X_2}} \quad \cdots \quad \underset{10\times 1}{\boxed{X_i}} \sim P_{data}(x)$$

$$D_{KL}(P_{data} \mid P_\theta) = E_{P_{data}}\left[\log \frac{P_{data}(x)}{P_\theta(x)}\right]$$

$$= E_{P_{data}}[\log P_{data}(x)] - E_{P_{data}}[\log P_\theta(x)]$$

$$= -\,\text{entropy}(P_{data}) - \underbrace{\boxed{E_{P_{data}}[\log P_\theta(x)]}}_{\max \theta}$$

$$\approx \underbrace{\frac{1}{n}\sum_{i=1}^{n} \log P_\theta(x_i)}_{\text{log likelihood}}$$

$$\log p(x) \approx \underbrace{\log p(x) - D_{KL}(\overbrace{q_p(z \mid x)}^{\text{tractable approximation}} \mid \overbrace{p(z \mid x)}^{\text{true posterior}})}_{\text{variational lower bound}}$$

$$q(z|x) \quad \begin{matrix} z \\ \downarrow \\ x \end{matrix} \quad = \log p(x) - E_{q(z|x)}\left(\log \frac{q(z|x)}{p(z|x)}\right)$$

$$= E_{q(z|x)}\left[\log p(x) - \log q(z|x) + \log p(z|x)\right]$$

$$= E_{q(z|x)}\left[-\log q(z|x) + \log p(x,z)\right] \qquad p(x,z) = p(z)p(x|z)$$

Bayesian network / Graphical model / Directed Acyclic Graph (DAG)

* similar to Asia example

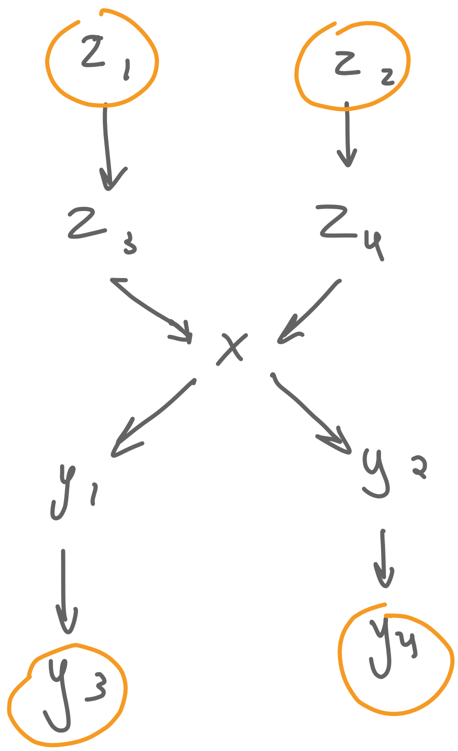

inference based on conditional dist.

$$p(x \mid z_1, z_2, y_3, y_4) = \frac{p(x, z_1, z_2, y_3, y_4)}{p(z_1, z_2, y_3, y_4)}$$

$\sum_{z_3 z_4 y_1 y_2}$

$\sum_{x, \dots}$

calculate using

$\longrightarrow$ get joint distribution $p(x, z, Y)$

$$= p(z_1)p(z_2)p(z_3|z_1)p(z_4|z_2)p(x|z_3,z_4)$$
$$p(y_1, y_2|x)p(y_3|y_1)\,p(y_4|y_2)$$

$$p(x \mid z_1, z_2, y_3, y_4) \propto \underbrace{p(x \mid z_1, z_2)}_{\text{prior for } x}\ \underbrace{p(y_3, y_4 \mid x)}_{\text{likelihood for } x}$$

normalization

$$\longrightarrow \frac{p(y_3 y_4\ x \mid z_1, z_2)}{p(y_3 y_4 \mid z_1, z_2)} = P(x \mid y_3 y_4, z_1, z_2)$$

marginalization

$$p(x \mid z_1, z_2) = \sum_{z_3 z_4} p(x\ z_3\ z_4 \mid z_1, z_2)$$

factorization

$$= \sum_{z_3 z_4} p(z_3 z_4 \mid z_1, z_2)\, p(x \mid z_3\ z_4\ z_1, z_2)$$

$$= \sum_{z_3 z_4} p(z_3 \mid z_1)\, p(z_4 \mid z_2)\, p(x \mid z_3, z_4)$$

$$P(y_3\, y_4 \mid x) = \underbrace{\sum_{y_1 y_2}}_{\text{marginalize}} P(y_1\, y_2\, y_3\, y_4 \mid x)$$

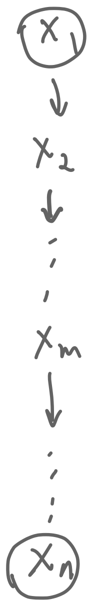

$$P(x_m \mid x_1, x_n) \propto \underbrace{p(x_m \mid x_1)}_{\substack{\pi(x_m) \\ \text{prior}}} \underbrace{P(x_n \mid x_m)}_{\substack{\lambda(x_m) \\ \text{likelihood}}}$$

belief propogation / message passing

$$P(x_{t+1} \mid x_1) = \sum_{x_t} P(x_t, x_{t+1} \mid x_1)$$

$$= \sum_{x_t} P(x_t \mid x_1)\, P(x_{t+1} \mid x_t)$$

$$\pi(x_{t+1}) = \sum_{x_t} \pi(x_t)\, P(x_{t+1} \mid x_t)$$

π message

Hidden Markov Model

$$\begin{tikzcd}
\text{hidden} & X_0 \arrow[r] & x_1 \arrow[r] \arrow[d] & x_2 \arrow[r] \arrow[d] & \cdots \arrow[r] & X_t \arrow[r, "\text{transition}"] \arrow[d, "\text{emission}"] & X_{t+1} \arrow[r] \arrow[d] & \cdots \\
\text{observed} & & y_1 & y_2 & & y_t & y_{t+1} &
\end{tikzcd}$$

$$P(X_t \mid y_1\, y_2 \ldots y_t) \longrightarrow P(X_{t+1} \mid y_1\, y_2 \ldots y_{t+1})$$

$$\overset{\shortparallel}{\tilde{p}(x_t)} \qquad \text{Forward algorithm} \qquad \overset{\shortparallel}{\tilde{P}(X_{t+1} \mid y_{t+1})}$$

$$\tilde{P}(\cdot) = P(\cdot \mid y_1 \ldots y_t)$$

$$\tilde{p}(x_t, x_{t+1}) = \tilde{P}(X_t) \cdot p(X_{t+1} \mid X_t)$$

$$\tilde{P}(X_{t+1}) = \sum \tilde{P}(x_t, x_{t+1})$$

$$\tilde{P}(X_{t+1} \mid y_{t+1}) \propto \tilde{P}(X_{t+1})\, P(Y_{t+1} \mid X_{t+1})$$

State space model

(X, Y) continuous

transition $[X_{t+1} \mid X_t] \sim N(AX, B)$

emission $[Y_t \mid X_t] \sim N(CX_t, D)$

$$\begin{tikzcd}
X_t \arrow[r, "{A,B}"] \arrow[d, "{C,D}"] & X_{t+1} \arrow[d] \\
Y_t & Y_{t+1}
\end{tikzcd}$$

Kalman Filtering

$$X_t = \begin{bmatrix} X_t^{(1)} \\ X_t^{(2)} \end{bmatrix} \begin{array}{l} \to \text{position} \\ \to \text{velocity} \end{array}$$

$Y_t \to$ noisy version of $X_t^{(1)}$

$$X_{t+1} = AX_t + B \qquad Y_t = C X_t + D$$

— Assuming constant velocity

$$\begin{bmatrix} X_{t+1}^{(1)} \\ X_{t+1}^{(2)} \end{bmatrix} = \overset{A}{\begin{bmatrix} 1 & \Delta t \\ 0 & 1 \end{bmatrix}} \begin{bmatrix} X_t^{(1)} \\ X_t^{(2)} \end{bmatrix} + B$$

$$X_t \xrightarrow{A,B} X_{t+1}$$

$$Y_t = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} X_t^{(1)} \\ X_t^{(2)} \end{bmatrix} + D$$

Assuming

$$\tilde{P}(X_t) \sim N(\mu, \Sigma), \text{ then}$$

$$\tilde{P}(X_{t+1}) \sim N(\underbrace{A\mu}_{\hat{\mu}}, \underbrace{A\Sigma A^T + B}_{\hat{\Sigma}})$$

To find $P(Y_{t+1} | X_{t+1})$

$$y_{t+1} = C X_{t+1} + \varepsilon_{t+1} \quad \text{where } \varepsilon_{t+1} \sim N(0, D)$$

and $\varepsilon_{t+1} \perp X_{t+1}$

Joint distribution of $X_{t+1} \sim N(\hat{\mu}, \hat{\Sigma})$ & $\varepsilon_{t+1} \sim N(0, D)$

$$\begin{bmatrix} X_{t+1} \\ \varepsilon_{t+1} \end{bmatrix} \sim N\left(\begin{bmatrix} \hat{\mu} \\ 0 \end{bmatrix}, \begin{bmatrix} \hat{\Sigma} & 0 \\ 0 & D \end{bmatrix} \right)$$

Joint distribution of X_{t+1} & Y_{t+1}

$$\begin{bmatrix} X_{t+1} \\ Y_{t+1} \end{bmatrix} = \begin{bmatrix} I & 0 \\ C & I \end{bmatrix} \begin{bmatrix} X_{t+1} \\ \varepsilon_{t+1} \end{bmatrix}$$

$$E\begin{bmatrix} \downarrow \\ \ \end{bmatrix} = \begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix} E \begin{bmatrix} X_{t+1} \\ \varepsilon_{t+1} \end{bmatrix} = \begin{bmatrix} \hat{\mu} \\ C\hat{\mu} \end{bmatrix}$$

$$\text{var}\begin{bmatrix} X_{t+1} \\ Y_{t+1} \end{bmatrix} = \begin{bmatrix} I & 0 \\ C & I \end{bmatrix} \text{Var}\begin{bmatrix} X_{t+1} \\ \varepsilon_{t+1} \end{bmatrix} \begin{bmatrix} I & C^T \\ 0 & I \end{bmatrix}$$

$$= \begin{bmatrix} I & 0 \\ C & I \end{bmatrix} \begin{bmatrix} \hat{\Sigma} & 0 \\ 0 & D \end{bmatrix} \begin{bmatrix} I & C^T \\ 0 & I \end{bmatrix} = ?$$

So $\begin{bmatrix} X_{t+1} \\ Y_{t+1} \end{bmatrix} \sim N\left(\begin{bmatrix} \hat{\mu} \\ C\hat{\mu} \end{bmatrix}, \quad ? \right)$

from joint dist, we can find

$(Y_{t+1} \mid X_{t+1}) \sim N(\mu_{new}, \Sigma_{new})$

where μ_{new} & Σ_{new} are kalman update based on new observation

Markov Decision Process

$$S_0 \to \cdots \to \underbrace{S_t \xrightarrow[\text{}]{\overset{\text{action}}{a_t}} S_{t+1}}_{r_t \text{ reward}} \to \cdots$$

state

(1) dynamics: $P(s_{t+1} \mid s_t, a_t)$

(2) policy : $P(a_t \mid s_t)$

(3) reward : $P(r_t \mid s_t, a_t, s_{t+1})$

return $R_t = r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \cdots$

γ ↓ discount factor

$\max \; E_{policy}(R_t)$

$\text{Value}^{policy}(s) = E_{policy}(R_t \mid S_t = s)$

$Q^{policy}(s, a) = E_{policy}(R_t \mid S_t = s, a_t = a)$