Bayesian decision theory

identification of the optimal decision to make, given a set of possible actions and uncertain beliefs encoded in $p(\theta|I)$ (the posterior of Bayesian inference, usually)

notations: $\{\theta\}$: set of features (= observable variable)
$\{a\}$: set of actions.

expected utility hypothesis: given a set of gain functions, the optimal decision rule is to take the action that maximizes the expected utility. $U(a|I)$ defined by:

$$U(a|I) = \langle G(a|\theta) \rangle_{p(\theta|I)}$$

$$\boxed{U(a|I) = \int G(a|\theta)\, p(\theta|I)\, d\theta}$$

$\Rightarrow$ take action $\boxed{a^* = \underset{a}{\arg\max}\; U(a|I)}$.

typical cases:

① summarizing an inference: we have inferred $p(\theta|d, X)$. we want to quote one value for θ. the groundtruth is $\tilde{\theta}$.

- zero-one gain functions: if $\tilde{\theta}$ is guessed correctly: 1, otherwise 0. valid for a discrete set of values for θ.
for a continuous problem the generalization is $G(\theta|\tilde{\theta}) = \delta_D(\theta-\tilde{\theta})$.

$$U(\theta|d,X) = \langle G(\theta|\tilde{\theta}) \rangle_{p(\tilde{\theta}|d,X)} = \int \delta_D(\theta-\tilde{\theta})\, p(\tilde{\theta}|d,X)\, d\tilde{\theta}$$

$$= p(\theta|d,X), \text{ to be maximized.}$$

$\Rightarrow$ the optimal decision is to quote the posterior mode:

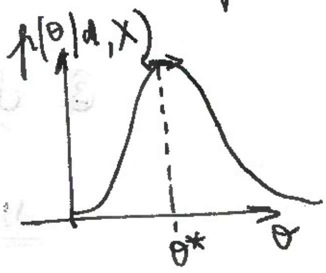

- minimizing the mean squared error: $G(\theta|\tilde{\theta}) = -(\tilde{\theta}-\theta)^2$.

$$U(\theta|d,X) = \langle G(\theta|\tilde{\theta})\rangle_{p(\tilde{\theta}|d,X)} = -\int (\theta-\tilde{\theta})^2 \, p(\tilde{\theta}|d,X)\, d\tilde{\theta}.$$

$\frac{\partial U}{\partial \theta} = 0$ yields the optimal value $\theta^* = \int \tilde{\theta}\, p(\tilde{\theta}|d,X)\, d\tilde{\theta} = \langle \theta \rangle$

i.e. the posterior mean. (can be different from the mode).

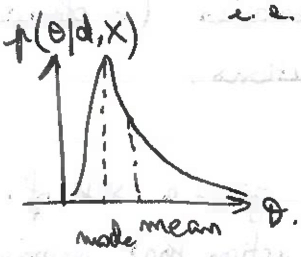

② Bayesian alerts : we look for an event E, we have access to $p(E|I)$. $p(\bar{E}|I) = 1 - p(E|I)$

two possible actions: a_1 = raise the alarm;
a_2 = do nothing.

$$U(a_1|I) = \underbrace{G(a_1|E)}_{\substack{\text{correct detection} \\ \text{(a "hit")}}} p(E|I) + \underbrace{G(a_1|\bar{E})}_{\substack{\text{false positive} \\ \text{(a "false alarm")}}} [1-p(E|I)].$$

$$U(a_2|I) = \underbrace{G(a_2|E)}_{\substack{\text{false negative} \\ \text{(a "miss")}}} p(E|I) + \underbrace{G(a_2|\bar{E})}_{\text{correct rejection}} [1-p(E|I)].$$

typical choice: (these can be actuals pounds!)

$$\begin{cases} G(a_1|E) = G - C \\ G(a_2|E) = 0 \end{cases} \qquad G(a_1|\bar{E}) = -C. \qquad G(a_2|\bar{E}) = 0.$$

(G: expected gain for a detection of E; C: cost of raising the alert)

$\Rightarrow U(a_1|I) = p(E|I)\,(G-C) + [1-p(E|I)]\,(-C)$ $\qquad U(a_2|I) = 0$

$U(a_1|I) \geq 0$? $\Rightarrow$ raise the alert if and only if $\boxed{p(E|I) \geq \frac{G}{C}}$

③ classifying patterns. see slides
+ Leclercq et al. 2015b, arXiv: 1503.00730

Notebook: DecisionTheory.