Decision Theory

Chris Williams

Probabilistic Modelling and Reasoning (INFR11134)
School of Informatics, The University of Edinburgh

Spring Semester 2024

Motivation

- Given tests **x** that have been made, the probability that a patient has cancer is determined as 0.4. Should their surgeon operate or not?
- You own a bakery, and have a probabilistic prediction for the demand for loaves tomorrow. How many loaves should you bake tomorrow morning?
- Decision theory tells us how to take actions rationally in the face of uncertainty

Outline

- Loss
- Risk
- Discrete-valued actions
- Real-valued actions

- Probabilistic inference allows us to update our beliefs about hidden quantities H given evidence $\mathbf{X} = \mathbf{x}$
- However, we often need to turn our beliefs into **actions**
- Q: How can we decide which action is best?
- A: Decision theory

Discussion largely based on Murphy PML1 (2022), sec 5.1

Loss

- An agent has a set of possible actions $\mathcal{A}$ to choose from
- E.g. a clinician may suspect that a patient has a tumour, and can either operate, or do nothing
- Operating on a patient if they don't have a tumour exposes them to possible harms (e.g., infections)
- Each action has costs and benefits, which depend on the underlying **state of nature** $h \in \mathcal{H}$
- **Loss function** $\ell(h, a)$ specifies the loss incurred when taking action a when the state of nature is h

Risk

- Given observations $\mathbf{x}$, we obtain $p(h|\mathbf{x})$
- Risk
$$R(a|\mathbf{x}) = \sum_h \ell(h, a)p(h|\mathbf{x})$$
(with sum replaced by an integral when necessary)
- Optimal policy
$$\pi^*(\mathbf{x}) = \operatorname*{argmin}_a R(a|\mathbf{x})$$

- Loss and utility: Basically same thing with opposite sign, $U(h, a) = -\ell(h, a)$
- Maximize expected utility, minimize expected loss

Example: should you cancel the concert?

- You are organizing an outdoor concert on Saturday
- On Thursday the weather forecast $\mathbf{x}$ for Saturday indicates a 60% chance of rain
- Losses are as follows

$$\ell(\text{fair}, \text{go ahead}) = -1 \qquad \ell(\text{rain}, \text{go ahead}) = 2$$

$$\ell(\text{fair}, \text{cancel}) = 3 \qquad \ell(\text{rain}, \text{cancel}) = 0$$

- Calculate the minimum risk strategy. Should you cancel the concert?

$$R(\text{go ahead}|\mathbf{x}) =$$

$$R(\text{cancel}|\mathbf{x}) =$$

Example: to operate or not?

- Patients have state of nature $h_1 = $ healthy, $h_2 = $ tumour.
- Actions are $a_1 = $ discharge the patient, $a_2 = $ operate
- Assume $\ell_{11} = \ell_{22} = 0$, $\ell_{12} = 1$ and $\ell_{21} = 10$, i.e. it is 10 times worse to discharge the patient when they have a tumour than to operate when they do not
- Risks

$$R(a_1|\mathbf{x}) = \ell_{11} P(h_1|\mathbf{x}) + \ell_{21} P(h_2|\mathbf{x}) = \ell_{21} P(h_2|\mathbf{x})$$
$$R(a_2|\mathbf{x}) = \ell_{12} P(h_1|\mathbf{x}) + \ell_{22} P(h_2|\mathbf{x}) = \ell_{12} P(h_1|\mathbf{x})$$

- Choose action a_1 when $R(a_1|\mathbf{x}) < R(a_2|\mathbf{x})$, i.e. when

$$\ell_{21} P(h_2|\mathbf{x}) < \ell_{12} P(h_1|\mathbf{x})$$

or

$$\frac{P(h_1|\mathbf{x})}{P(h_2|\mathbf{x})} > \frac{\ell_{21}}{\ell_{12}} = 10$$

- If $\ell_{21} = \ell_{12}$ then threshold is 1; in our case we require stronger evidence in favour of h_1 (healthy) in order to discharge the patient

Real-valued actions

- L2 loss $\ell(h, a) = (h - a)^2$
- Risk

$$R(a|\mathbf{x}) = \mathbb{E}[(h-a)^2|\mathbf{x}] = \mathbb{E}[h^2|\mathbf{x}] - 2a\,\mathbb{E}[h|\mathbf{x}] + a^2$$

- Minimize wrt a

$$\frac{\partial}{\partial a} R(a|\mathbf{x}) = -2\mathbb{E}[h|\mathbf{x}] + 2a = 0$$

- Solution

$$\pi^*(\mathbf{x}) = \mathbb{E}[h|\mathbf{x}]$$

  i.e., to pick the **posterior mean**
- For L1 loss $\ell(h, a) = |h - a|$, solution is the **posterior median**

The Utility of Money

- A utility curve assigns numeric values to various possible outcomes
- In general, most people's utility function tends to be concave for positive amounts of money
- Thus the incremental value of additional money decreases as wealth grows
- Empirical psychological studies show that people's utility functions often grow logarithmically in the amount of monetary gain

Source: Koller and Friedman (2009), sec 22.2.1

Summary

- Optimal action minimizes the risk (expected loss)
- In general losses are cost-sensitive $\ell(h, a)$
- To minimize L2 loss, predict the posterior mean
- Not covered: sequential decision problems