Practice problems! Decision problems and estimation

Christos Dimitrakakis: chrdimi@chalmers.se

Time and scoring The numbers indicate expected time to complete the exercise. If you take more time than indicated, please note that. A "!" indicates that it may require some extra thought. A "?" indicates that this is an open question. The exercises count towards your grade.

1 Problems with no observations (30-45 min)

For the first part of exercises, we consider a set of worlds Ω and a decision set $\mathcal{D}$, as well as the following utility function $U : \Omega \times \mathcal{D} \to \mathbb{R}$:

$$U(\omega, a) = \operatorname{sinc}(\omega - d) \tag{1.1}$$

where $\operatorname{sinc}(x) = \sin(x)/x$. If ω is known and $\mathcal{D} = \Omega = \mathbb{R}$ then obviously the optimal decision is $d = \omega$, as $\operatorname{sinc}(x) \leq \operatorname{sinc}(0) = 1$. However, we consider the following case:

$$\Omega = \mathcal{D} = \{-2.5, \ldots, -0.5, 0, 0.5, \ldots, 2.5\}.$$

Exercise 1 (15)**.** *Assume ω is drawn from ξ, with $\xi(\omega) = 1/11$ for all $\omega \in \Omega$, calculate and plot the expected utility $U(\xi, d) = \sum_\omega \xi(\omega) U(\omega, d)$ for each d. Report* $\max_d U(\xi, d)$

Exercise 2 (5)**.** *Assume $\omega \in \Omega$ is arbitrary (but deterministically selected). Calculate the utility $U(d) = \min_\omega U(\omega, d)$ for each d. Report* $\max(U)$.

Exercise 3 (10 (+ 40!))**.** *Again assume $\omega \in \Omega$ is arbitrary (but deterministically selected). We now allow for stochastic policies π on D. Then the expected utility is $U(\omega, \pi) = \sum_d U(\omega, d)\pi(d)$.*

(a) (10) Calculate and plot the expected utility when $\pi(d) = 1/11$ for all d, reporting values for all ω.

(b) (40!) Bonus question: Find

$$\max_\pi \min_\xi U(\xi, \pi).$$

Hint: Use the linear programming formulation, adding a constant to the utility matrix U so that all elements are non-negative.

2 Problems with observations (60-90 min)

For this section, we consider a set of worlds Ω and a decision set D, as well as the following utility function $U : \Omega \times D \to \mathbb{R}$:

$$U(\omega, d) = -|\omega - d|^2. \tag{2.1}$$

In addition, we consider a family of distributions on a sample space $S = \{0, 1\}^n$,

$$\mathscr{F} \triangleq \{f_\omega \mid \omega \in \Omega\}, \tag{2.2}$$

such that f_ω is the binomial probability mass function with parameters ω (with the number of draws n being implied).

Consider the parameter set:

$$\Omega = \{0, 0.1, \ldots, 0.9, 1\}. \tag{2.3}$$

Let ξ be the uniform distribution on Ω, such that $\xi(\omega) = 1/11$ for all $\omega \in \Omega$. Let the decision set be:

$$D = [0, 1]. \tag{2.4}$$

Exercise 4 (2)**.** *What is the decision d^* maximising $U(\xi, d) = \sum_\omega \xi(\omega)U(\omega, d)$ and what is $U(\xi, d^*)$?*

Exercise 5 (40)**.** *In the same setting, we now observe the sequence $x = (x_1, x_2, x_3) = (1, 0, 1)$.*

1. *(5) Plot the posterior distribution $\xi(\omega \mid x)$ and compare it with the posterior we would obtain if our prior on ω was $\xi' = \mathcal{B}eta(2, 2)$.*
2. *(10) Find the decision d^* maximising the* a posteriori *expected utility*

$$\mathbb{E}_\xi(U \mid d, x) = \sum_\omega U(\omega, d)\xi(\omega \mid x).$$

3. *(25) Consider $n = 2$, i.e. $S = \{0, 1\}^2$. Calculate the Bayes-optimal expected utility in extensive form:*

$$\mathbb{E}_\xi(U \mid \delta^*) = \sum_S \phi(x) \sum_\omega U[\omega, \delta^*(x)]\xi(\omega \mid x) = \sum_S \phi(x) \max_d \sum_\omega U[\omega, d]\xi(\omega \mid x), \quad (2.5)$$

   where $\phi(x) = \sum_\omega f_\omega(x)\xi(\omega)$ is the prior marginal distribution of x and $\delta^* : S \to D$ is the Bayes-optimal decision rule.

   Hint: You can simplify the computational complexity somewhat, since you only really need to calculate the probability of $\sum_t x_t$. This is not necessary to solve the problem though.

Exercise 6 (20)**.** *In the same setting, we consider nature to be adversarial. Once more, we observe $x = (1, 0, 1)$. Assume that nature can choose a prior among a set of priors $\Xi = \{\xi_1, \xi_2\}$. Let $\xi_1(\omega) = 1/11$ $\xi_2(\omega) = \omega/5.5$.*

1. *Calculate and plot our value for deterministic decisions d:*

$$\min_{\xi \in \Xi} \mathbb{E}_\xi(U \mid d, x).$$

2. *Find the minimax prior ξ^**

$$\min_{\xi \in \Xi} \max_{\delta \in \mathscr{D}} \mathbb{E}_\xi(U \mid \delta)$$

Hint: Apart from the adversarial prior selection, this is very similar to the previous exercise.