Lecture Notes for CS497 Jan 28 and 30

Hanning Zhou
hzhou@ifp.uiuc.edu
University of Illinois at Urbana-Champaign, 405 N. Mathews, Urbana, IL 61801

1 Lecture 1/28/03

Another example of decision theory: parameter estimation
Let χ denote a parameter space
$\chi = \Theta = U$ $\chi \subseteq Y$ (others are possible)
Nature picks $x \in \chi$. Try to guess it.

Suppose the observation y has normal distribution with mean x, i.e. $p(y|x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-y)^2}{2\sigma^2}}$ and x has prior $x \sim Unif[0,1]$ i.e. a prior $p(x) = \{ \begin{array}{ll} 1 & x \in [0,1] \\ 0 & otherwise \end{array}$; $\chi = [0,1]$. The loss function measures estimation error $L(u,x) = (u-x)^2$

First find posterior $p(x|y) = \frac{p(y|x)p(x)}{p(y)}$, and associated risk $C(u|y) = \int_\chi L(u,x)p(x|y)dx$.

Then select parameter $u = argminC(u|y)$, since $p(y)$ is common over all u, so it can be omitted, i.e. $u = argmin \int_\chi L(u,x)p(y|x)p(x)dx$

In the case of multiple observations $C(u|y) = \int_\chi L(u,x)p(y_1|x)\cdots p(y_k|x)p(x)dx$.

2 Utility Theory

Where does L come from ? R: reward $= -L$
Example:
u_1: shovelling manure u_2: do nothing
Case 1: $R(u_1) = \$1000$ $R(u_2) = \$0$
Most people will prefer u_1
Case 2: $R(u_1) = \$10,001,000$ $R(u_2) = \$10,000,000$
Most people will prefer u_2, this shows that utility of unit money is decreasing when the total amount increases as shown in Figure 2.
Case 3: $R(u_1)$:\$10,000
$R(u_2)$: toss a coin with probability $\frac{1}{2}$ to get \$25,000
Most people (except a few gambler) choose u_1.
Case 4: Play 100 times
$R(u_1)$:\$100
$R(u_2)$:get \$250 with probability $\frac{1}{2}$
$\Re$: reward space. Take case 3 for instance, $\Re = \{10000, 25000, 0\}$

Suppose

	u_1	u_2
θ_1	1	0
θ_2	3	2

$p(\theta_1) = \frac{1}{4}$ $p(\theta_2) = \frac{3}{4}$
If choose u_1: $r = 1$ with probability $\frac{1}{4}$.
If choose u_1: $r = 1$ with probability $\frac{3}{4}$.
p: a probability distribution over $\Re$. $p \in P$.
State preferences between any p_1, $p_2 \in P$. Either $p_1 \prec p_2$ (p_2 is preferred over p_1) or $p_2 \prec p_1$ or $p_1 \simeq p_2$.

Goal of utility theory: find some function $V : \Re \rightarrow R$ that represents preference in P, such that $p_1 \prec p_2$ iff $E^{p_1}[V(r)] < E^{p_2}[V(r)]$

When does such V exist? Only when the DM is rational.

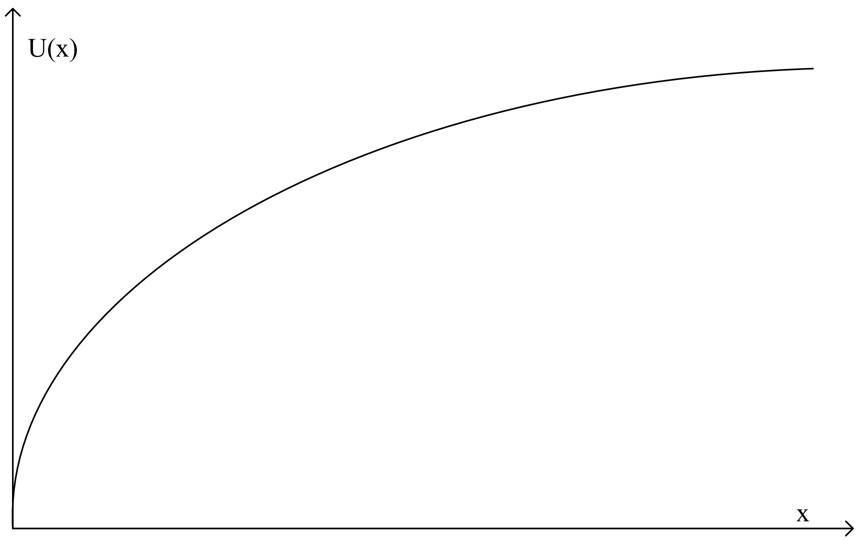

Figure 1: The utility R(x) versus the amount of money x.

1. If $p_1, p_2 \in P$, then $p_1 \preceq p_2$ or $p_2 \preceq p_1$: "You must be able to make up your mind"
2. If $p_1 \preceq p_2$ and $p_2 \preceq p_3$ then $p_1 \preceq p_3$. Transitivity of preferences: This is not always reasonable (counter e.g. Rock Paper Scissor)
3. If $p_1 \prec p_2$ then $\alpha p_1 + (1-\alpha)p_3 \prec \alpha p_2 + (1-\alpha)p_3$ $\forall \alpha \in (0,1) \forall p_3 \in P$ "Don t be confused by new lotteries"
4. If $p_1 \prec p_2 \prec p_3$ then $\exists \alpha \in (0,1); \beta \in (0,1)$ s.t. $\alpha p_1 + (1-\alpha)p_3 \preceq p_2$ (p_3 is not heaven); $p_2 \preceq \beta p_1 + (1-\beta)p_3$ (p_1 is not hell). "No heaven or hell"

Criticism of decision theory
1. Values seem subjective
2. Difficult to assign losses
3. Sloppy losses $\rightarrow$ poor decisions
-Sensitivity analysis helps sometimes

3 Lecture 1/30/03

3.1 Nondeterministic decision making

Suppose we use rewards

	u_1	u_2
θ_1	2	1000
θ_2	2	0

Take a check for \$1000 or take \$2 in cash.

Minimax choice u_1, receive \$2. We pay a price to have guaranteed earnings!
Backward justification using Bayesian analysis, take the priors: $p(\theta_1) \ll p(\theta_2)$
Bayesian decision making, what are probability?

Bayesian	frequentist
prob. is belief about one trial	prob. is only what is obtained in the limit after many repeated trials
subjective, personal	objective, minimalist

Frequentist risk: $R(\theta, \gamma) = \int_Y L(\theta, \gamma(y)) p(y|\theta) dy$ $\gamma : Y \rightarrow U$
What are prior distributions? $p(\theta)$
Laplace:"principle of insufficient reason" 1812
All natural outcomes are equally likely
Construct noninformative priors: Shannon entropy $H(p) = -\sum_{\Theta} p(\theta) \log(p(\theta)) d\theta$
$H(p) = 0$ for uniform $p(\theta)$
Criticism
1. What if $\Theta = R$? -Use generalized pdf's.
2. e.g. θ_1: no precipitation; θ_2: precipitation
Or θ_1: no precipitation; θ_2: rain; θ_3: snow

Selection of the set of θ is arbitrary and informative.

Example: Fitting observation of 2-dimensional points (x_i, y_i) with a straight line. $\theta = (\theta_1, \theta_2, \theta_3)$; $\theta_1 x + \theta_2 y + \theta_3 = 0$ and $\theta_1^2 + \theta_2^2 + \theta_3^2 = 1$. The distribution of $\theta_1, \theta_2, \theta_3$ can be

1. Uniform in unit area on a sphere in 3 dimensional space of $\theta_1, \theta_2, \theta_3$.
2. Uniform in longitude and latitude on the sphere.

Parameterizing of θ is arbitrary.

3.2 Multi-objective decision making

No uncertainties. Scenario

1. U; 2. $L : U \rightarrow R^d$; 3. Find all $u \in U$ s.t. $Not \exists u' \in U$ s.t. $L(u') \preceq L(u)$

What is $\preceq$? For $L(u) = [L_1(u), L_2(u) \dots L_d(u)]$

$L(u) \preceq L(u')$ iff $L_i \leq L'_i$ $\forall i \in \{1, 2, \cdots, d\}$

Example 1: Two robots A and B are scheduled to move from their origins to the destinations in the shortest time, as shown in Figure 3.2. The solution is $[t_A t_B]$.

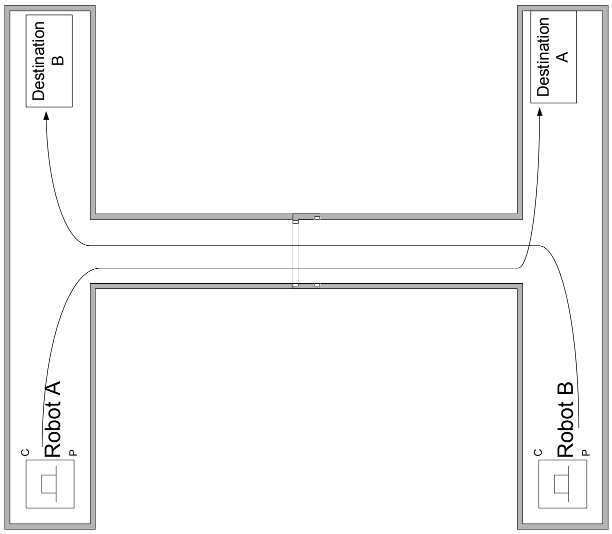

Figure 2: Robots scheduling

u_1: Robot A goes first, robot B waits, then goes as soon as possible. $L(u_1) = [35]$

u_2: Robot B goes first, robot A waits, then goes as soon as possible. $L(u_2) = [53]$

u_3: Both robots go until they detect each other, negotiate and backs up. $L(u_3) = [64]$

u_1 and u_2 are Pareto optimal.

(Counter) Example 2: $U = [0, 1] \subseteq R$; $L(u) = [u, 1 - u]$. Can't order any of the pairs (u, u').