Lectures Notes #1 Jan 21

Decision Theory: 1 DM (Decision Maker), 1 Stage
Sequential Decision Making: 1 DM, n Stages
Game Theory: m DMs, 1 Stage

m DMs, n Stages
Alternative Names for decision makers: robot, agent, player, bot.

1 DM, 1 Stage Single Stage Decision Making
Scenario 0
1. Let U be a set of choices $\{u_1, u_2, u_3, \ldots, u_n\}$
2. Let $L : U \to \Re$ be the ("loss", "cost") function
3. Select $u \in U$ that minimizes the loss $L(u)$

U can be equivalently called "choices", "actions", or "inputs".
L can be equivalently called "loss", "cost", or "negative reward".

Reward and the Loss function are related: $R(u) = -L(u)$
We will want to select some u that maximizes R, or that equivalently minimizes L.

Randomized (mixed) vs. Deterministic (Pure) decisisions
Deterministic: select one u_i
Randomized: select u_i with probability p_i $\forall u_i \in U$. The probability distribution on U satisfies the obvious:

$$p_i \geq 0$$

$$\sum_{i=1}^{n} p_i = 1 \ \forall i \in \{1, \ldots, n\}$$

.

Let u^* denote a randomized strategy.
Let U denote the set of all randomized strategies. This is nothing but a bunch of vectors of the form

$$[p_1, p_2, \ldots, p_n]$$

where we assign some probability for each p_i.
Note that we can have a deterministic strategy by simply making one of the $p_i = 1$ and the rest of them zero.

Senario 1
1. U as before.
2. L as before.
3. Select u^* from U that minimizes the expected loss

$$E(L) = \sum_{i=1}^{n} L(u_i) \cdot p_i$$

Note that random is good only if we can play the game over and over again.

Example: Which is more helpful Scenario 1 or Scenario 2?
Consider the following loss function

$$L = [2\ 3\ 5]$$

best choice is the first one, since it has the smallest loss. So let $p_1 = 1$ and $p_2 = p_3 = 0$
Now if we had instead

$$L = [2\ 3\ 2]$$

then the best choice would be either the first one or the last one.
In this case a Deterministic choice will always yield an optimal solution, and the Randomized strategy is worse!! However Scenario 1 is not quite useless.

Expected value is good for Game Theory, so that you are not predictable to your enimies (Lot harder to guess my action then).

Matching Pennies
Player 1 chooses H or T
Player 2 chooses H or T (simultaneously)
One player doesn't have knowledge of what the other chose. If outcome is HH or TT, P1 payes P2 \$1. Otherwise P2 pays P1 \$1.

If you have a deterministic policy here then your opponent might learn your strategy. So Randomized is better. But it might take a long time to learn if P1 is deterministic or not.
For example consider the digits of π as controlling the decisions of P1. It would appear to P2 as if P1 is random.

Now suppose that we have a continuoum of choices
Scenario 0b
1. $U \subseteq \Re^d$ (d is the dimension)
2. $L : U \rightarrow \Re$
3. Select $u \in U$ to minimize L.

Example, let $L(u) = u^2$, $u \in [-1,\ 1]$. Since we are trying to minimize L, then L has its minimum obviously at $u = 0$. Now let's say that $u \in (0,\ 1)$. Since we can't really choose zero, then their is no true answer (theoritically) better then any other. Offcourse we can say that the closest number to zero is whatever IEEE says that floating point number should be for a particular system.
This problem can also arise when a robot is trying to find the shortest path from point A to point B and there are obstucles in its path. Then the question again becomes how close can we get to the obstucles.

There is a fix to this problem:
Define inf = greatest lower bound (similar to min)
sup = lowest upper bound (similar to max)
These two definitions allow us to select the number that we always wanted to (in our case this would be zero if we apply the inf to the set $(0,\ 1)$).

Reasoning about UNCERTAINTY: a game against nature

Nature is a special DM (decision maker). We don't know what decision nature will make (or has made).
Let θ be a choice for nature. $\theta \in \Theta$, where Θ is the set of all choices for nature.

$$\Theta = \{\theta_1, \theta_2, \ldots, \theta_m\} \textit{ the Discrete version}$$

or

$$\Theta \subseteq \Re^d \textit{ the continuous version.}$$

The Loss function is then

$$L : U \times \theta \rightarrow \Re$$

For example, consider the following table

$$
\begin{bmatrix}
 & u_1 & u_2 & u_3 \\
\theta_1 & 1 & -1 & 2 \\
\theta_2 & -1 & 2 & -1 \\
\theta_3 & 0 & -2 & 1
\end{bmatrix}
$$

Two ways to make a decision

Nondeterministic: I have no idea - use Murphy's Law to make a decision (assume the worst case scenario). In our example we would choose u_1.

Probabilistic: I've been observing nature and gathering statistics.