# CS497 Week 1 Notes (1/21, 1/23)

Scribe: Steve Lindemann

# 1 Single-Stage Decision Making

## 1.1 Introductory concepts

To introduce some of the basic concepts in single-stage decision making, consider the following scenario:

**Scenario 0**

1. *Let $U$ be a set of possible choices:* $\{u_1, u_2, \ldots, u_n\}$.
2. *Let $L : U \to \mathbb{R}$ be a* loss function *or* cost function.
3. *Select a $u \in U$ that minimizes $L(U)$.*

In this scenario, we see that the set $U$ consists of all choices that we can make; these are also called *actions* or *inputs*. The loss function $L$ represents the cost associated with each possible choice; another approach is to define a *reward function* $R$ which represents the gain or benefit of each choice. These approaches are equivalent, since on can simply take $R(u) = -L(u)$.

A method used to make a decision is called a *strategy*. In this scenario, our strategy was *deterministic*; that is, given some set $U$ and function $L$, our choice is completely determined. Alternatively, we could have taken a *randomized* strategy, in which our decision also depended on the outcome of some random events. In this strategy, we define a function $p : U \to \mathbb{R}$ such that the probability of selecting a particular choice $u$ is $p(u)$; denote $p(u_i) = p_i$. The ordinary rules governing probability spaces apply (e.g., $\sum_{i=1}^{n} p_i = 1, p_i \geq 0 \ \forall i$). Randomized and deterministic strategies are also called *mixed* and *pure*, respectively. For purposes of notation, we will use $u^*$ to refer to a randomized strategy and $\mathcal{U}$ to refer to the set of all randomized strategies.

**Example 1** *Let the input set $U = \{a, b\}$. Then one can choose a randomized strategy $u^*$ in the following way:*

1. *Flip a fair H/T coin.*
2. *If the result is H, choose a; if T, choose b.*

*Since the coin is fair, this corresponds to choosing $p(a) = 0.5, p(b) = 0.5$.*

Consider the following scenario:

**Scenario 1**

1. $U = \{u_1, u_2, \ldots, u_n\}$
2. $L : U \to \mathbb{R}$

3. *Select* $u^* \in \mathcal{U}$ *that minimizes* $\mathrm{E}[L] = \sum_{i=1}^{n} L(u_i)p_i$.

$\mathrm{E}[L]$ reflects the average loss if the game were to be played many times. Now, Scenarios 0 and 1 are identical, with the exception that one uses a deterministic strategy, and one uses a randomized strategy. Which is better? To help answer this, we give the following example:

**Example 2** *Let* $U = \{1, 2, 3\}$, *and* $L(1) = 2, L(2) = 3, L(3) = 5$ *(we may write this in vector notation as* $L = [2\ 3\ 5]$*). Following the deterministic strategy from Scenario 0, we choose* $u = 1$*. What if we use the strategy from Scenario 1? By inspection we can see that we need* $p = [1\ 0\ 0]$*; thus, the randomized strategy results in the same choice as the deterministic one.*

We have seen in the above example that a randomized strategies and deterministic ones can produce identical results. However, what if for some input set $U$ and loss function $L$, we have $L(u_i) = L(u_j)$? Then, there can be randomized strategies which act differently than deterministic ones. However, if one only considers the minimum loss attained, they are not better because both types of strategies will select actions resulting in minimum loss. Thus, in this case we find that Scenario 1 is useless! However, randomized strategies are very useful in general, as shown in the following example.

**Example 3 (Matching Pennies)** *Consider a game in which two players simultaneously choose H or T. If the outcome is HH or TT (the players choose the same), then Player 1 pays Player 2 \$1; if the outcome is HT or TH, then Player 2 pays Player 1 \$1. What happens if Player 1 uses a deterministic strategy? If Player 2 can determine what that strategy is, then he can choose his strategy so that he always wins the game. However, if Player 1 chooses a randomized strategy, he can at least expect to break even (what randomized strategy guarantees this?).*

So far, we have examined scenarios in which there were only a finite number of possible choices. Many problems, however, have a continuum of choices, as does the following:

**Scenario 2**

1. $U \subseteq \mathbb{R}^d$ *(usually,* $U$ *is closed and bounded)*
2. $L : U \to \mathbb{R}$
3. *Select* $u \in U$ *to minimize* $L$

This is a classical optimization problem.

**Example 4** *Let* $U = [-1, 1] \subset \mathbb{R}$ *and* $L(u) = u^2$*. To attain minimum cost we choose* $u = 0$*.*

However, what if in the example above we chose $U = (0, 1)$? Then the minimum is not well-defined. However, we can introduce the concept of the *infimum*, which is the greatest lower bound of a set. Similarly, we can introduce the *supremum*, which is the least upper bound of a set. Then, we can still say $\inf_{u \in U} L(u) = 0$.

## 1.2 Reasoning about uncertainty: a game against nature

In the previous scenarios, we have assumed complete knowledge about the loss function $L$. This need not be the case, however; in particular situations, there may be uncertainty involved. One convenient way to describe this uncertainty is to introduce a special decision-maker, called *nature*. Nature is an unreasoning entity (i.e., it is not an adversary), and we do not know what decision nature will make (or has made). We call the set $\Theta$ the set of choices for nature (alternatively, the *parameter space*), and $\theta \in \Theta$ is a particular choice by nature. The parameter space may be either discrete or continuous; in the discrete case, we have $\Theta = \{\theta_1, \theta_2, \ldots, \theta_n\}$, and in the continuous case we have $\Theta \subseteq \mathbb{R}^d$. Then, we can define the loss function to be $L : U \times \Theta \to \mathbb{R}$, in which the operator $\cdot \times \cdot$ is the Cartesian product.

**Example 5** *Let L be specified by the following table:*

$U$

| $\Theta$ | | | |
|---|---|---|---|
| | 1 | −1 | 2 |
| | −1 | 2 | −1 |
| | 0 | −2 | 1 |

The best strategy to adopt depends on what model we have of what nature will do:

- **Nondeterministic**: I have no idea.
- **Probabilistic**: I have been observing nature and gathering statistics.

In the first case, one might assume Murphy's Law ("If anything can go wrong, it will"); then, one would choose the column with the least maximum value. Alternatively, one might assume that nature's decisions follow a uniform distribution, all choices being equally likely. Then one would choose the column with the least average loss (this approach was taken by Laplace in 1812). In the second case, one could use Bayesian analysis to calculate a probability distribution $P(\theta)$ of the actions of nature, and use that to make decisions. The following two scenarios formalize these approaches.

**Scenario 3 (Minimax solution)**

1. $U = \{u_1, \ldots, u_n\}$
2. $\Theta = \{\theta_1, \ldots, \theta_m\}$
3. $L : U \times \Theta \to \mathbb{R}$
4. *Choose* $u$ *to minimize* $\max_{\theta \in \Theta} L(u, \theta)$.

**Scenario 4 (Expected optimal solution)**

1. $U = \{u_1, \ldots, u_n\}$
2. $\Theta = \{\theta_1, \ldots, \theta_m\}$
3. $P(\theta)$ *given* $\forall \theta \in \Theta$
4. $L : U \times \Theta \to \mathbb{R}$
5. *Choose* $u$ *to minimize* $\mathrm{E}^{\theta}[L] = \sum_{\theta \in \Theta} L(u, \theta) P(\theta)$.

Again consider Example 5. If the strategy from Scenario 3 is adopted, then we would choose $u_1$ so that we would pay loss 1 in the worst case. If the strategy from Scenario 4 is chosen, and assuming $P(\theta_1) = 1/5$, $P(\theta_2) = 1/5$, $P(\theta_3) = 3/5$, we find that $u_2$ has the lowest expected loss, and so would take that action. If the probability distribution had been $P = [1/10\ 4/5\ 1/10]$, then simple calculations show that $u_1$ is the best choice. Hence our decision depends on $P(\theta)$; if this information is statistically valid, then better decisions are made. If it is not, then potentially worse decisions can be made.

Another strategy is to minimize "regret", the amount of loss you could have eliminated if you had chosen differently, given the action of nature. A regret matrix corresponding to Example 5 can be found in Figure 1. Given some regret matrix, one can adopt a minimax or expected optimal strategy.

| | $U$ | | |
|---|---|---|---|
| | 2 | 0 | 3 |
| $\Theta$ | 0 | 3 | 0 |
| | 2 | 0 | 3 |

Figure 1: A regret matrix corresponging to Example 5.

## 1.3 Having a single observation

Let $y$ be an *observation*; this could be some data, a measurement, or a sensor reading. Let $Y$ be the *observation space*, the set of all possible $y$. Now, we can make a decision based on $y$; let $\gamma : Y \to U$ denote a *decision rule* (strategy, plan). Then modify our decision strategies as follows:

- **Nondeterministic**: Assume there is some $F(y) \subseteq \Theta$, which is known for every $y \in Y$. Choose some $\gamma$ that minimizes $\max\limits_{\theta \in F(y)} L(\gamma(y), \theta)$ for each $y \in Y$.

- **Probabilistic**: Assume that $P(y|\theta)$ is known, $\forall y \in Y, \forall \theta \in \Theta$. Then Bayes rule yields $P(\theta|y) = P(y|\theta)P(\theta)/P(y)$, in which $P(y) = \sum_{\theta\in\Theta} P(y|\theta)P(\theta)$[1] Then choose $\gamma$ so that it minimizes the *conditional Bayes risk* $R(u|y) = \sum_{\theta\in\Theta} L(u, \theta)P(\theta|y)$, for every $y \in Y$.

Formally, we have the following scenarios:

**Scenario 5 (Nondeterministic)**

1. *$U = \{u_1, \ldots, u_n\}$*
2. *$\Theta = \{\theta_1, \ldots, \theta_m\}$*
3. *$Y = \{y_1, \ldots, y_l\}$*
4. *$F(y)$ given $\forall y \in Y$*
5. *$L : U \times \Theta \to \mathbb{R}$*
6. *Choose $\gamma$ to minimize $\max\limits_{\theta \in F(y)} L(\gamma(y), \theta)$ for each $y \in Y$.*

**Scenario 6 (Bayesian decision theory)**

1. *$U = \{u_1, \ldots, u_n\}$*
2. *$\Theta = \{\theta_1, \ldots, \theta_m\}$*
3. *$Y = \{y_1, \ldots, y_l\}$*
4. *$P(\theta)$ given $\forall \theta \in \Theta$.*
5. *$P(y|\theta)$ given $\forall y \in Y, \theta \in \Theta$*
6. *$L : U \times \Theta \to \mathbb{R}$*

[1]For the purposes of decision-making, $P(y)$ is simply a scaling factor and may be omitted.

7. *Choose $\gamma$ to minimize $R(\gamma(y)|y)$ for every $y \in Y$.*

Extending the former case, we may imagine that we have $k$ observations: $y_1, \ldots, y_k$. Then, $R(u|y_1, \ldots, y_k) = \sum_{\theta \in \Theta} L(u, \theta) P(\theta|y_1, \ldots, y_k)$. If we assume that $P(y_i|\theta)$ is known for each $i \in \{1, \ldots, k\}$ and that conditional independence holds, we have $P(\theta|y_1, \ldots, y_k) = \left(\prod_{i=1}^{k} P(y_i|\theta)\right) P(\theta)/P(y_1, \ldots, y_k)$.

**Example 6 (Classification)** *Let $\Omega = \{\omega_1, \ldots, \omega_n\}$ denote a set of* classes*, and let $y$ denote a* feature *and $Y$ a* feature space*. For this type of problem, we have $\Theta = U = \Omega$, since nature selects an object from one of the classes, and we attempt to identify the class nature has selected. The feature set $Y$ represents useful information that can help us identify which class an object belongs to. A reasonable cost function is*

$$L(u, \theta) = \begin{cases} 0 & \text{if } u = \theta \text{ (the classification is correct)} \\ 1 & \text{if } u \neq \theta \text{ (the classification is incorrect)} \end{cases}$$

*If the Bayesian decision strategy is adopted, it will result in choices that minimize the expected probability of misclassification.*

**Example 7 (Optical Character Recognition)** *Let $\Omega = \{A, B, C, D, E, F, G, H\}$. Further, imagine that we our image processing algorithms can extract the following features:*

| | | |
|---|---|---|
| *Shape* | *0* | *A E F H* |
| | *1* | *B C D G* |
| *Ends* | *0* | *B D* |
| | *1* | |
| | *2* | *A C G* |
| | *3* | *F E* |
| | *4* | *H* |
| *Holes* | *0* | *C E F G H* |
| | *1* | *A D* |
| | *2* | *B* |

*Assuming that the image processing algorithms never err, we can use a minimax strategy to make our decision. Are there any letters which the features do not distinguish? If so, what enhancements might we make to our image processing algorithms to distinguish them? If we assume that that the image processing algorithms sometimes make mistakes, then we can use a Bayesian strategy. After running the algorithms thousands of times and gathering statistics, we can learn the necessary conditional probabilities and use them to make the decision with the highest expectation of success.*