10-716: Advanced Machine Learning **Spring 2019**

Lecture 4: January 24

Lecturer: Pradeep Ravikumar *Scribes: Menglan Ji, Terrance Liu, Elan Rosenfeld*

4.1 Minimax Analysis

Minimax analysis can be used to choose actions. In this model, we have two players, nature and the ML person. Natures chooses an action first, the ML person second. However, this does not mean the ML person gets to see the choice of nature.

$$\overline{V} = \inf_{a \in \mathcal{A}} \sup_{\theta \in \Theta} L(\theta, a) \geq \sup_{\theta \in \Theta} \inf_{a \in \mathcal{A}} L(\theta, a) = \underline{V} \tag{4.1}$$

In strategy $\overline{V}$, nature first chooses θ and then the ML person chooses a, without seeing θ. It is a conservative choice of action. Because the ML person doesn't know nature's choice, it must look at the worst case choice of nature. In other words, θ is not fixed. This is why $\overline{V} \geq \underline{V}$.

While Equation 4.1 is about deterministic choices, it is also applicable to randomized choices.

$$\sup_{\theta \in \Theta} L(\theta, a) = \sup_{\pi \in \Pi} L(\pi, a)$$

where Π is the set of distributions over θ. If we fix a, then the supremum is a single θ, and the distribution is a point mass at θ. If there are multiple θs, then the distribution has mass at all θs.

$$\inf_{\delta^*} L(\theta, \delta^*) = \inf_{a} L(\theta, a)$$

for a particular choice of θ. Using randomized choices doesn't change the value of the game.

The **minimax strategy** is $\delta^m \in \arg\inf_{\delta^*} \sup_{\theta} L(\theta, \delta^*)$. This is actionable, so we can solve the problem, but it is not as easy as Bayesian estimators, which focus on each x individually.

The **maximin strategy** is $\pi^m \in \arg\sup_{\pi \in \Pi} L(\pi, \delta^*)$. Here, nature maximizes the loss incurred.

Theorem 4.1 *Suppose π_0^*, δ_0^* such that $L(\theta, \delta_0^*) \leq L(\pi_0^*, a)$ $\forall \theta \in \Theta, a \in \mathcal{A}$. Then, δ_0^* is minimax for the ML person, and π_0^* is maximin for nature.*

Proof:

$$\overline{V} = \inf_{\delta^*} \sup_{\pi \in \Pi} L(\pi, \delta^*) \leq \sup_{\pi \in \Pi} L(\pi, \delta_0^*) \leq \inf_{a \in \mathcal{A}} L(\pi_0^*, a) \leq \sup_{\pi \in \Pi} \inf_{a \in \mathcal{A}} L(\pi_0^*, a) = \underline{V}$$

Because $\underline{V} \leq \overline{V}$, the inequalities are all equalities. Thus, δ_0^* is the supremum and π_0^* is the infimum. ■

This is not constructive, but can be used to verify that an estimator is minimax.

4.2 Equalizer Rules

One method for determining a minimax decision rule is to search for an equalizer rule.

Theorem 4.2 *If δ_0^* is Bayes with respect to prior π over Θ, and π is least favorable with respect to δ_0^*, i.e. $R(\theta, \delta_0^*) \le r(\pi, \delta_0^*)$, then δ_0^* is minimax.*

Proof:

$$\begin{aligned}
\overline{V} &= \inf_{\delta^*} \sup_{\theta} R(\theta, \delta^*) \\
&\le \sup_{\theta} R(\theta, \delta_0^*) \\
&\le r(\pi, \delta_0^*) && \text{upper bounded by Bayes risk} \\
&= \inf_{\delta} \gamma(\pi_0, \delta) && \text{precondition that } \delta_0^* \text{ is Bayes wrt } \pi \\
&\le \sup_{\pi} \inf_{\delta} r(\pi_0, \delta) && \pi_0 \text{ is least favorable prior} \\
&\le \underline{V}
\end{aligned}$$

■

Theorem 4.3 *Suppose δ_0^* is a decision rule s.t.*

1. $R(\theta, \delta_0^*) = C \ \forall \theta \in \Theta$
2. δ_0^* *is Bayes w.r.t. some prior* π *over* Θ

Then δ_0^* is minimax.

Proof: Let us choose π^* s.t. δ_0^* is Bayes w.r.t π^*. Then the Bayes risk is given by

$$\begin{aligned}
r(\pi^*, \delta_0^*) &= E_{\pi^* \sim \Theta}[R(\theta, \delta_0^*)] \\
&= E_{\pi^* \sim \Theta}[C] \\
&= C
\end{aligned}$$

Therefore, we have that $R(\theta, \delta_0^*) = C = r(\pi^*, \delta_0^*) \implies R(\theta, \delta_0^*) \le r(\pi^*, \delta_0^*)$, i.e. π^* is a least favorable prior. Using Theorem 4.2, we conclude that δ_0^* is minimax.

■

In some cases, focusing solely on the worst-case risk is not appropriate.

Example Consider Figure 4.1, which plots the risk $R(\theta, \delta_i)$ against values of θ. In plot on the left, minimax analysis would prefer the line in red, given that the worst-case risk is lower. Now suppose for certain values for θ, we increase $R(\theta, \delta_i)$ by some constant. Intuitively, it would not make sense for our decision rule preferences to change if the information changed is common to both choices of δ_i. However, as shown on the plot on the right, the decision rule denoted by the red plot now has a larger worst-case risk.

Example Suppose the table of possible states of nature θ and actions a, with their respective losses $L(\theta, a)$, are as follows:

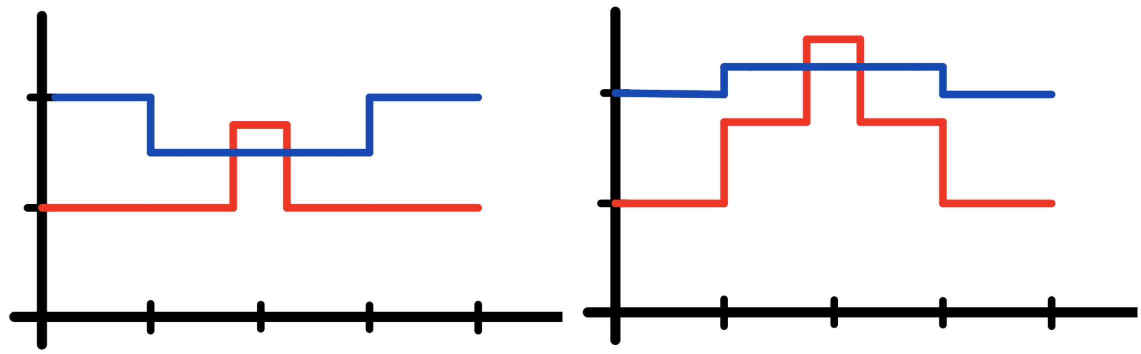

Figure 4.1: Minimax analysis

	a_1	a_2
θ_1	10	10.01
θ_2	8	-8

Looking at the worst cases

$$\sup_\theta L(\theta, a_2) = 10.01$$
$$\sup_\theta L(\theta, a_1) = 10$$

so we would pick a_1. However

$$L(\theta_1, a_1) \approx L(\theta_1, a_2)$$
$$L(\theta_2, a_1) >> L(\theta_2, a_2)$$

so we would actually prefer a_2. However, the minimax rule would choose a_1.

Example Suppose $\theta \in (0, 1), \mathcal{A} = [0, 1]$. Think it as a biased coin, flipped n times, so $X \sim Bin(n, \theta)$. We will use the scaled loss capped at 2, $L(\theta, a) = \min\{\frac{(\theta-a)^2}{\theta^2}, 2\}$.

The minimax rule is $\delta^m(x) = 0$, because it achieves the lowest risk.

If we look at $\delta_0(x) = 0$, we can see that $L(\theta, \delta_0(x)) = 1$, and we take any other $\delta \neq \delta_0$. Let

$$c = \min_{x \in B_\delta} \delta(x) > 0$$

where $B_\delta = \{x \,:\, \delta(x) \neq 0\}$ is nonempty because $\delta \neq \delta_0$. We have $L(\theta, \delta(x)) = 2 \; \forall \theta < \frac{c}{1+\sqrt{2}}$. The reason is that we are measuring relative error, and $\delta(x)$ does not take values smaller than c. For $\theta < \frac{c}{1+\sqrt{2}}$, we have

$$\begin{aligned}
R(\theta, \delta) &= \mathbb{E}_X L(\theta, \delta(x)) \\
&= \sum_{x \in B_\delta} L(\theta, \delta(x)) + \sum_{x \notin B_\delta} L(\theta, \delta(x)) \\
&= \sum_{x \in B_\delta} 2f(x \mid \theta) + \sum_{x \notin B_\delta} f(x \mid \theta) \\
&= \sum_{x} f(x \mid \theta) + \sum_{x \in B_\delta} f(x \mid \theta) > 1 \ni L(\theta, \delta_0(x)) \approx R(\theta, \delta_0)
\end{aligned}$$

So we have $1 = \sup_\theta R(\theta, \delta_0) < \sup_\theta R(\theta, \delta) \; \forall \delta \neq \delta_0$.

4.3 Minimax Regret

Definition 4.4 *Minimax regret is the worst case difference between the loss from the action chosen under δ and the minimum possible loss under that same supremizing θ.*

$$R_{minimax} = \inf_{\delta} \sup_{\theta} \left\{ L(\theta, a) - \inf_{a'} L(\theta, a') \right\}$$

Example Suppose the table of possible states of nature θ and actions a, with their respective losses $L(\theta, a)$, are as follows:

	a_1	a_2
θ_1	10	10.01
θ_2	8	-8

If the state of nature is θ_1, then a_2 is only slightly worse than a_1, and if instead the state of nature is θ_2, then a_2 is significantly better. Thus, a_2 appears to be the better action. However, optimizing for minimax risk would suggest picking a_1, because the maximum loss is smaller.

If instead we consider minimax regret, we arrive at a more reasonable decision. The table corresponding to $L(\theta, a) - \inf_{a'} L(\theta, a')$ is

	a_1	a_2
θ_1	0	0.01
θ_2	16	0

Because θ is supremized in $R_{minimax}$, we pick the action with the *smallest worst-case* regret. For a_1 the worst-case regret is 16 and for a_2 it is 0.01. Therefore we pick a_2.

Recall that the risk of a policy δ is

$$r(\pi, \delta) = \mathbb{E}_\theta \mathbb{E}_x \left[L(\theta, \delta(x)) \right]$$

The expected regret for a policy δ can be expressed as

$$\mathbb{E}_\theta \mathbb{E}_x \left[L(\theta, \delta(x)) - \inf_a L(\theta, a) \right] = \underbrace{\mathbb{E}_\theta \mathbb{E}_x \left[L(\theta, \delta(x)) \right]}_{r(\pi,\delta)} - \mathbb{E}_{\theta \sim \pi} \left[\inf_a L(\theta, a) \right] \tag{4.2}$$

The second term on the right hand side of Equation 4.2 does not depend on δ at all. So we have:

$$r^*(\pi, \delta) = r(\pi, \delta) - \mathbb{E}_{\theta \sim \pi} \left[\inf_a L(\theta, a) \right] \tag{4.3}$$

Thus, whether you minimize the Bayes risk, or whether you minimize the Bayes regret, you will find the same decision rule. The minimax rule and the minimax regret rule are going to be different, however.

4.4 Loss Functions

Returning to the consideration of loss functions, let's start with a look at the squared loss: $L(\theta, a) = (\theta - a)^2$.

There's a sense in which this can be seen as an approximation to a function of the difference between the two terms. Suppose we have $L(\theta, a) = g(\theta - a)$, where $g''(0) > 0$. We can write the Taylor Expansion of this as

$$g(\theta - a) \approx g(0) + g'(0)(\theta - a) + \frac{g''(0)}{2}(\theta - a)^2 \approx (\theta - a + c)^2$$

Observe that this is similar to squared loss, with a difference of some constant c.

There are several other popular loss functions. One is the scaled absolute loss, defined as

$$L(\theta, a) = \begin{cases} k_1(\theta - a), & \theta \geq a \\ k_2(a - \theta), & \theta < a \end{cases}$$

If $k_1 = k_2 = k$ then this can just be written as $L(\theta, a) = k|\theta - a|$. We can also consider the zero-one loss, defined as

$$L(\theta, a) = \mathbb{I}\{\theta \neq a\}$$

4.5 Hypothesis Testing

Recall that in hypothesis testing, we are considering two possible (sets of) states of nature and we wish to take an action based on which we think is true. Under the null $\theta \in \Theta_0$ we may wish to take action a_0 (e.g., retain the null) and under the alternative $\theta \in \Theta_1$ we prefer action a_1 (reject the null).

Suppose we decide on the zero-one loss, such that $L(\theta, a) = \mathbb{I}\{\theta \notin \Theta_i\}$.

The risk of a policy δ under a given state of nature θ is therefore

$$R(\theta, \delta) = \mathbb{E}_x \left[\mathbb{I}\{\theta \notin \Theta_{\delta(x)}\}\right]$$

We can see that under the null,

$$R(\theta_0, \delta) = \mathbb{P}_{x \sim p(\cdot|\theta_0)}(\delta(x) \neq a_0) = \underbrace{\mathbb{P}_{x \sim p(\cdot|\theta_0)}(\delta_x \text{ is incorrect})}_{\text{Type I Error}}$$

Similarly,

$$R(\theta_1, \delta) = \underbrace{\mathbb{P}_{x \sim p(\cdot|\theta_1)}(\delta_x \text{ is incorrect})}_{\text{Type II Error}}$$

So we can see that the two types of error in hypothesis testing can be reframed as the risk of a policy under the null or alternative hypotheses using zero-one loss.