

Lecture Notes 14

36-705

We continue with our discussion of decision theory.

1 Decision Theory

Suppose we want to estimate a parameter θ using data $X^n = (X_1, \dots, X_n)$. What is the best possible estimator $\hat{\theta} = \hat{\theta}(X_1, \dots, X_n)$ of θ ? Decision theory provides a framework for answering this question.

1.1 The Risk Function

Let $\hat{\theta} = \hat{\theta}(X^n)$ be an estimator for the parameter $\theta \in \Theta$. We start with a **loss function** $L(\theta, \hat{\theta})$ that measures how good the estimator is. For example:

$$\begin{aligned} L(\theta, \hat{\theta}) &= (\theta - \hat{\theta})^2 && \text{squared error loss,} \\ L(\theta, \hat{\theta}) &= |\theta - \hat{\theta}| && \text{absolute error loss,} \\ L(\theta, \hat{\theta}) &= |\theta - \hat{\theta}|^p && L_p \text{ loss,} \\ L(\theta, \hat{\theta}) &= 0 \text{ if } \theta = \hat{\theta} \text{ or } 1 \text{ if } \theta \neq \hat{\theta} && \text{zero-one loss,} \\ L(\theta, \hat{\theta}) &= I(|\hat{\theta} - \theta| > c) && \text{large deviation loss,} \\ L(\theta, \hat{\theta}) &= \int \log \left(\frac{p(x; \theta)}{p(x; \hat{\theta})} \right) p(x; \theta) dx && \text{Kullback-Leibler loss.} \end{aligned}$$

If $\theta = (\theta_1, \dots, \theta_k)$ is a vector then some common loss functions are

$$\begin{aligned} L(\theta, \hat{\theta}) &= \|\theta - \hat{\theta}\|^2 = \sum_{j=1}^k (\hat{\theta}_j - \theta_j)^2, \\ L(\theta, \hat{\theta}) &= \|\theta - \hat{\theta}\|_p = \left(\sum_{j=1}^k |\hat{\theta}_j - \theta_j|^p \right)^{1/p}. \end{aligned}$$

When the problem is to predict a $Y \in \{0, 1\}$ based on some classifier $h(x)$ a commonly used loss is

$$L(Y, h(X)) = I(Y \neq h(X)).$$

For real valued prediction a common loss function is

$$L(Y, \hat{Y}) = (Y - \hat{Y})^2.$$




Figure 1: Comparing two risk functions. Neither risk function dominates the other at all values of θ .

The **risk** of an estimator $\hat{\theta}$ is

$$R(\theta, \hat{\theta}) = \mathbb{E}_{\theta} \left(L(\theta, \hat{\theta}) \right) = \int L(\theta, \hat{\theta}(x_1, \dots, x_n)) p(x_1, \dots, x_n; \theta) dx. \quad (1)$$

When the loss function is squared error, the risk is just the MSE (mean squared error):

$$R(\theta, \hat{\theta}) = \mathbb{E}_{\theta} (\hat{\theta} - \theta)^2 = \text{Var}_{\theta}(\hat{\theta}) + \text{bias}^2. \quad (2)$$

If we do not state what loss function we are using, assume the loss function is squared error.

1.2 Comparing Risk Functions

To compare two estimators, we compare their risk functions. However, this does not provide a clear answer as to which estimator is better. Consider the following examples.

Example 1 Let $X \sim N(\theta, 1)$ and assume we are using squared error loss. Consider two estimators: $\hat{\theta}_1 = X$ and $\hat{\theta}_2 = 3$. The risk functions are $R(\theta, \hat{\theta}_1) = \mathbb{E}_{\theta} (X - \theta)^2 = 1$ and $R(\theta, \hat{\theta}_2) = \mathbb{E}_{\theta} (3 - \theta)^2 = (3 - \theta)^2$. If $2 < \theta < 4$ then $R(\theta, \hat{\theta}_2) < R(\theta, \hat{\theta}_1)$, otherwise, $R(\theta, \hat{\theta}_1) < R(\theta, \hat{\theta}_2)$. Neither estimator uniformly dominates the other; see Figure 1.

Example 2 Let $X_1, \dots, X_n \sim \text{Bernoulli}(p)$. Consider squared error loss and let $\hat{p}_1 = \bar{X}$. Since this has zero bias, we have that

$$R(p, \hat{p}_1) = \text{Var}(\bar{X}) = \frac{p(1-p)}{n}.$$

Another estimator is

$$\hat{p}_2 = \frac{Y + \alpha}{\alpha + \beta + n}$$

where $Y = \sum_{i=1}^n X_i$ and α and β are positive constants.¹ Now,

$$\begin{aligned} R(p, \hat{p}_2) &= \text{Var}_p(\hat{p}_2) + (\text{bias}_p(\hat{p}_2))^2 \\ &= \text{Var}_p\left(\frac{Y + \alpha}{\alpha + \beta + n}\right) + \left(\mathbb{E}_p\left(\frac{Y + \alpha}{\alpha + \beta + n}\right) - p\right)^2 \\ &= \frac{np(1-p)}{(\alpha + \beta + n)^2} + \left(\frac{np + \alpha}{\alpha + \beta + n} - p\right)^2. \end{aligned}$$

Let $\alpha = \beta = \sqrt{n/4}$. The resulting estimator is

$$\hat{p}_2 = \frac{Y + \sqrt{n/4}}{n + \sqrt{n}}$$

and the risk function is

$$R(p, \hat{p}_2) = \frac{n}{4(n + \sqrt{n})^2}.$$

The risk functions are plotted in Figure 2. As we can see, neither estimator uniformly dominates the other.

These examples highlight the need to be able to compare risk functions. To do so, we need a one-number summary of the risk function. Two such summaries are the maximum risk and the Bayes risk.

The **maximum risk** is

$$\bar{R}(\hat{\theta}) = \sup_{\theta \in \Theta} R(\theta, \hat{\theta}) \quad (3)$$

and the **Bayes risk** under prior π is

$$B_\pi(\hat{\theta}) = \int R(\theta, \hat{\theta})\pi(\theta)d\theta. \quad (4)$$

Example 3 Consider again the two estimators in Example 2. We have

$$\bar{R}(\hat{p}_1) = \max_{0 \leq p \leq 1} \frac{p(1-p)}{n} = \frac{1}{4n}$$

¹This is the posterior mean using a Beta (α, β) prior.




Figure 2: Risk functions for \hat{p}_1 and \hat{p}_2 in Example 2. The solid curve is $R(\hat{p}_1)$. The dotted line is $R(\hat{p}_2)$.

and

$$\bar{R}(\hat{p}_2) = \max_p \frac{n}{4(n + \sqrt{n})^2} = \frac{n}{4(n + \sqrt{n})^2}.$$

Based on maximum risk, \hat{p}_2 is a better estimator since $\bar{R}(\hat{p}_2) < \bar{R}(\hat{p}_1)$. However, when n is large, $\bar{R}(\hat{p}_1)$ has smaller risk except for a small region in the parameter space near $p = 1/2$. Thus, many people prefer \hat{p}_1 to \hat{p}_2 . This illustrates that one-number summaries like the maximum risk are imperfect.

These two summaries of the risk function suggest two different methods for devising estimators: choosing $\hat{\theta}$ to minimize the maximum risk leads to minimax estimators; choosing $\hat{\theta}$ to minimize the Bayes risk leads to Bayes estimators.

An estimator $\hat{\theta}$ that minimizes the Bayes risk is called a **Bayes estimator**. That is,

$$B_\pi(\hat{\theta}) = \inf_{\tilde{\theta}} B_\pi(\tilde{\theta}) \quad (5)$$

where the infimum is over all estimators $\tilde{\theta}$. An estimator that minimizes the maximum risk is called a **minimax estimator**. That is,

$$\sup_{\theta} R(\theta, \hat{\theta}) = \inf_{\tilde{\theta}} \sup_{\theta} R(\theta, \tilde{\theta}) \quad (6)$$

where the infimum is over all estimators $\tilde{\theta}$. We call the right hand side of (6), namely,

$$R_n \equiv R_n(\Theta) = \inf_{\tilde{\theta}} \sup_{\theta \in \Theta} R(\theta, \tilde{\theta}), \quad (7)$$

the **minimax risk**. Statistical decision theory has two main goals: determine the minimax risk R_n and find an estimator that achieves this risk.

Once we have found the minimax risk R_n we want to find the minimax estimator that achieves this risk:

$$\sup_{\theta \in \Theta} R(\theta, \hat{\theta}) = \inf_{\hat{\theta}} \sup_{\theta \in \Theta} R(\theta, \hat{\theta}). \quad (8)$$

1.3 Bayes Estimators

Let π be a prior distribution. After observing $X^n = (X_1, \dots, X_n)$, the posterior distribution is, according to Bayes' theorem,

$$\mathbb{P}(\theta \in A | X^n) = \frac{\int_A p(X_1, \dots, X_n | \theta) \pi(\theta) d\theta}{\int_{\Theta} p(X_1, \dots, X_n | \theta) \pi(\theta) d\theta} = \frac{\int_A \mathcal{L}(\theta) \pi(\theta) d\theta}{\int_{\Theta} \mathcal{L}(\theta) \pi(\theta) d\theta} \quad (9)$$

where $\mathcal{L}(\theta) = p(x^n; \theta)$ is the likelihood function. The posterior has density

$$\pi(\theta | x^n) = \frac{p(x^n | \theta) \pi(\theta)}{m(x^n)} \quad (10)$$

where $m(x^n) = \int p(x^n | \theta) \pi(\theta) d\theta$ is the **marginal distribution** of X^n . Define the **posterior risk** of an estimator $\hat{\theta}(x^n)$ by

$$r(\hat{\theta} | x^n) = \int L(\theta, \hat{\theta}(x^n)) \pi(\theta | x^n) d\theta. \quad (11)$$

Theorem 4 *The Bayes risk $B_\pi(\hat{\theta})$ satisfies*

$$B_\pi(\hat{\theta}) = \int r(\hat{\theta} | x^n) m(x^n) dx^n. \quad (12)$$

Let $\hat{\theta}(x^n)$ be the value of θ that minimizes $r(\hat{\theta} | x^n)$. Then $\hat{\theta}$ is the Bayes estimator.

Proof:

Let $p(x, \theta) = p(x | \theta) \pi(\theta)$ denote the joint density of X and θ . We can rewrite the Bayes risk as follows:

$$\begin{aligned} B_\pi(\hat{\theta}) &= \int R(\theta, \hat{\theta}) \pi(\theta) d\theta = \int \left(\int L(\theta, \hat{\theta}(x^n)) p(x | \theta) dx^n \right) \pi(\theta) d\theta \\ &= \int \int L(\theta, \hat{\theta}(x^n)) p(x, \theta) dx^n d\theta = \int \int L(\theta, \hat{\theta}(x^n)) \pi(\theta | x^n) m(x^n) dx^n d\theta \\ &= \int \left(\int L(\theta, \hat{\theta}(x^n)) \pi(\theta | x^n) d\theta \right) m(x^n) dx^n = \int r(\hat{\theta} | x^n) m(x^n) dx^n. \end{aligned}$$

If we choose $\hat{\theta}(x^n)$ to be the value of θ that minimizes $r(\hat{\theta}|x^n)$ then we will minimize the integrand at every x and thus minimize the integral $\int r(\hat{\theta}|x^n)m(x^n)dx^n$.

Now we can find an explicit formula for the Bayes estimator for some specific loss functions.

Theorem 5 *If $L(\theta, \hat{\theta}) = (\theta - \hat{\theta})^2$ then the Bayes estimator is*

$$\hat{\theta}(x^n) = \int \theta \pi(\theta|x^n) d\theta = \mathbb{E}(\theta|X = x^n). \quad (13)$$

If $L(\theta, \hat{\theta}) = |\theta - \hat{\theta}|$ then the Bayes estimator is the median of the posterior $\pi(\theta|x^n)$. If $L(\theta, \hat{\theta})$ is zero-one loss, then the Bayes estimator is the mode of the posterior $\pi(\theta|x^n)$.

Proof:

We will prove the theorem for squared error loss. The Bayes estimator $\hat{\theta}(x^n)$ minimizes $r(\hat{\theta}|x^n) = \int (\theta - \hat{\theta}(x^n))^2 \pi(\theta|x^n) d\theta$. Taking the derivative of $r(\hat{\theta}|x^n)$ with respect to $\hat{\theta}(x^n)$ and setting it equal to zero yields the equation $2 \int (\theta - \hat{\theta}(x^n)) \pi(\theta|x^n) d\theta = 0$. Solving for $\hat{\theta}(x^n)$ we get 13.

Example 6 *Let $X_1, \dots, X_n \sim N(\mu, \sigma^2)$ where σ^2 is known. Suppose we use a $N(a, b^2)$ prior for μ . The Bayes estimator with respect to squared error loss is the posterior mean, which is*

$$\hat{\theta}(X_1, \dots, X_n) = \frac{b^2}{b^2 + \frac{\sigma^2}{n}} \bar{X} + \frac{\frac{\sigma^2}{n}}{b^2 + \frac{\sigma^2}{n}} a. \quad (14)$$

It is worth keeping in mind the trade-off: Bayes estimators although easy to compute are very subjective; they depend strongly on the prior π . Minimax estimators, although more challenging to compute are not subjective, but do have the drawback that they are protecting against the worst-case which might lead to pessimistic conclusions.

2 Minimax Estimators through Bayes Estimators

Our goal is to compute a minimax estimator $\hat{\theta}$ that satisfies:

$$\sup_{\theta \in \Theta} R(\theta, \hat{\theta}) \leq \inf_{\tilde{\theta}} \sup_{\theta \in \Theta} R(\theta, \tilde{\theta}).$$

We will let θ_{minimax} denote a minimax estimator.

2.1 Bounding the Minimax Risk

One strategy to find the minimax estimator is by finding (upper and lower) bounds on the minimax risk that match. Then the estimator that achieves the upper bound is a minimax estimator.

Upper bounding the minimax risk is straightforward. Given an estimator $\hat{\theta}_{\text{up}}$ we can compute its maximum risk and use it to upper bound the minimax risk, i.e.

$$\inf_{\tilde{\theta}} \sup_{\theta \in \Theta} R(\theta, \tilde{\theta}) \leq R(\theta, \hat{\theta}_{\text{up}}).$$

The Bayes risk of the Bayes estimator for any prior π lower bounds the minimax risk. Fix a prior π and suppose that $\hat{\theta}_{\text{low}}$ is the Bayes estimator with respect to π , then we have that:

$$B_{\pi}(\hat{\theta}_{\text{low}}) \leq B_{\pi}(\theta_{\text{minimax}}) \leq \sup_{\theta} R(\theta, \theta_{\text{minimax}}) = \inf_{\tilde{\theta}} \sup_{\theta \in \Theta} R(\theta, \tilde{\theta}).$$

Let us see an example of this in action.

Example: We will prove a classical result that if we observe independent draws from a d -dimensional Gaussian, $X_1, \dots, X_n \sim N(\theta, I_d)$, then the average:

$$\hat{\theta} = \frac{1}{n} \sum_{i=1}^n X_i,$$

is a minimax estimator of θ with respect to the squared loss.

Let R_n denote the minimax risk. First, let us compute the upper bound on R_n . We note that,

$$\hat{\theta} \sim N(\theta, I_d/n),$$

so that its risk:

$$R(\theta, \hat{\theta}) = \mathbb{E}\left[\sum_{i=1}^d (\hat{\theta}_i - \theta_i)^2\right] = \mathbb{E}\left[\sum_{i=1}^d Z_i^2\right],$$

where $Z_i \sim N(0, 1/n)$. This yields that,

$$\inf_{\tilde{\theta}} \sup_{\theta \in \Theta} R(\theta, \tilde{\theta}) \leq R(\theta, \hat{\theta}) = \frac{d}{n}.$$

Now we lower bound the minimax risk using the Bayes risk. Let us take the prior to be zero-mean Gaussian, i.e. we take $\pi = N(0, c^2 I_d)$. By sufficiency, we can replace the data with $\hat{\theta}$. We can write:

$$\begin{aligned} \theta &\sim N(0, c^2 I_d) \\ \hat{\theta}|\theta &\sim N(\theta, I_d/n). \end{aligned}$$

We can write this as

$$\begin{aligned}\theta &= c\epsilon \\ \hat{\theta} &= \frac{1}{\sqrt{n}} Z\end{aligned}$$

where $\epsilon, Z \sim N(0, I_d)$. Hence,

$$\begin{pmatrix} \theta \\ \hat{\theta} \end{pmatrix} \sim N \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{bmatrix} c^2 I_d & c^2 I_d \\ c^2 I_d & (c^2 + 1/n) I_d \end{bmatrix} \right]$$

We can now compute the posterior (using standard conditional Gaussian formulae), and obtain its mean:

$$\mathbb{E}[\theta|\hat{\theta}] = \frac{c^2}{c^2 + 1/n} \hat{\theta}.$$

Now,

$$R(\theta, \hat{\theta}) = \mathbb{E} \left\| \frac{c^2}{c^2 + 1/n} \hat{\theta} - \theta \right\|^2.$$

Write $\hat{\theta} = \theta + W$, where $W \sim N(0, I_d/n)$. Then

$$R(\theta, \hat{\theta}) = \mathbb{E}_W \left\| \frac{c^2}{c^2 + 1/n} Z - \frac{\theta}{n(c^2 + 1/n)} \right\|^2.$$

Let us denote $\beta := c^2 + 1/n$. Then we obtain that,

$$R(\theta, \hat{\theta}) = \frac{\|\theta\|_2^2}{n^2 \beta^2} + \frac{c^4}{\beta^2} \mathbb{E} \|W\|_2^2 = \frac{\|\theta\|_2^2}{n^2 \beta^2} + \frac{c^4}{\beta^2} \frac{d}{n}.$$

The Bayes risk further averages this over $\theta \sim N(0, c^2 I_d)$ to obtain that,

$$B_\pi \left(\frac{c^2}{c^2 + 1/n} \hat{\theta} \right) = \frac{c^2 d}{n^2 \beta^2} + \frac{c^4}{\beta^2} \frac{d}{n} = \frac{c^2 d}{n \beta} = \frac{d}{n(1 + 1/(nc^2))}.$$

We conclude that

$$\frac{d}{n(1 + 1/(nc^2))} \leq R_n \leq \frac{d}{n}.$$

This is true for every $c > 0$. Since c was arbitrary we can take the limit as $c \rightarrow \infty$ to obtain that the minimax risk is upper and lower bounded by d/n and hence, $R_n = d/n$ and the sample average $\hat{\theta}$ is minimax.

2.2 Least Favorable Prior

The other way to obtain Bayes estimators is by constructing what are called least favorable priors.

Theorem 7 *Let $\hat{\theta}$ be the Bayes estimator for some prior π . If*

$$R(\theta, \hat{\theta}) \leq B_\pi(\hat{\theta}) \text{ for all } \theta \quad (15)$$

*then $\hat{\theta}$ is minimax and π is called a **least favorable prior**.*

Proof:

Suppose that $\hat{\theta}$ is not minimax. Then there is another estimator $\hat{\theta}_0$ such that $\sup_\theta R(\theta, \hat{\theta}_0) < \sup_\theta R(\theta, \hat{\theta})$. Since the average of a function is always less than or equal to its maximum, we have that $B_\pi(\hat{\theta}_0) \leq \sup_\theta R(\theta, \hat{\theta}_0)$. Hence,

$$B_\pi(\hat{\theta}_0) \leq \sup_\theta R(\theta, \hat{\theta}_0) < \sup_\theta R(\theta, \hat{\theta}) \leq B_\pi(\hat{\theta}) \quad (16)$$

which is a contradiction.

Theorem 8 *Suppose that $\hat{\theta}$ is the Bayes estimator with respect to some prior π . If the risk is constant then $\hat{\theta}$ is minimax.*

Proof:

The Bayes risk is $B_\pi(\hat{\theta}) = \int R(\theta, \hat{\theta})\pi(\theta)d\theta = c$ and hence $R(\theta, \hat{\theta}) \leq B_\pi(\hat{\theta})$ for all θ . Now apply the previous theorem.

Example 9 *Consider the Bernoulli model with squared error loss. We showed previously that the estimator*

$$\hat{p} = \frac{\sum_{i=1}^n X_i + \sqrt{n/4}}{n + \sqrt{n}}$$

has a constant risk function. This estimator is the posterior mean, and hence the Bayes estimator, for the prior $\text{Beta}(\alpha, \beta)$ with $\alpha = \beta = \sqrt{n/4}$. Hence, by the previous theorem, this estimator is minimax.