

Lecture 10: Le Cam’s Method (Some Examples)

Lecturer: Kirthevasan Kandasamy

Scribed by: Ying Fu, Deep Patel

Disclaimer: *These notes have not been subjected to the usual scrutiny reserved for formal publications. They may be distributed outside this class only with the permission of the instructor.*

Previously, we defined the notion of **Minimax Optimality**, following which we obtained lower bounds for this **Minimax Risk** by ‘reducing’ the problem of **estimation** to that of **hypothesis testing**. Upon this ‘reduction,’ we then looked at **Le Cam’s** method to obtain a **lower bound** by considering binary hypothesis testing specifically. In this lecture, we begin by obtaining a specific form of lower bound as a consequence of Le Cam’s method and then see its applications through some examples of mean estimation and regression.

Corollary 1 (From Le Cam’s Method). *Let $P_0, P_1 \in \mathcal{P}$ and $\delta = \rho(\theta(P_0), \theta(P_1))$. If $\text{KL}(P_0, P_1) \leq \frac{1}{n} \log 2$. Suppose the dataset S has n i.i.d samples. Then,*

$$R_n^* \geq \frac{1}{8} \Phi \left(\frac{\delta}{2} \right) \quad (1)$$

Proof From our last lecture, we know that,

$$\begin{aligned} R_n^* &\geq \frac{1}{2} \Phi \left(\frac{\delta}{2} \right) \|P_0^n \wedge P_1^n\| \\ &\geq \frac{1}{2} \Phi \left(\frac{\delta}{2} \right) \frac{1}{2} \exp(-\text{KL}(P_0^n, P_1^n)) \quad (\text{By the inequality : } \|P_0 \wedge P_1\| \geq \frac{1}{2} \exp(-\text{KL}(P_0, P_1))) \\ &= \frac{1}{4} \Phi \left(\frac{\delta}{2} \right) \exp(-n\text{KL}(P_0, P_1)) \quad (\text{Since } \text{KL}(P_0^n, P_1^n) = n\text{KL}(P_0, P_1)) \\ &\geq \frac{1}{8} \Phi \left(\frac{\delta}{2} \right) \quad (\text{By assumption}) \end{aligned}$$

□

Remark 0.1. *Corollary 1 formalizes the intuition that if statistically speaking, we can find two distributions that are close to each other (i.e. low KL-divergence) but their parameters, $\theta(P_0)$ and $\theta(P_1)$ are far apart (i.e. high δ), then the minimax optimal error is also high. This is because the lower bound, as seen above in the Corollary, will be high owing to difficulty in distinguishing between the two distributions via the binary hypothesis test.*

1 Estimating Mean of Normal Distribution (with known variance)

Suppose we have $S = \{X_1, \dots, X_n\}$ with n i.i.d samples drawn from $P \in \mathcal{P}$, where $\mathcal{P} = \{\mathcal{N}(\mu, \sigma^2) \mid \mu \in \mathbb{R}\}$ with σ^2 is known. We also have $\Phi \circ \rho(\theta_1, \theta_2) = (\theta_1 - \theta_2)^2$. Define $P_0 = \mathcal{N}(0, \sigma^2)$ and $P_1 = \mathcal{N}(\delta, \sigma^2)$ such

that $|\theta(P_0) - \theta(P_1)| = \delta$. Then the KL divergence between P_0 and P_1 is,

$$\text{KL}(P_0, P_1) = \frac{\delta^2}{2\sigma^2} \quad \left(\text{Since } \text{KL}(\mathcal{N}(\mu_1, \sigma^2), \mathcal{N}(\mu_2, \sigma^2)) = \frac{1}{2\sigma^2}(\mu_1 - \mu_2)^2 \right)$$

In order to satisfy the requirement of Corollary 1 – which is $\text{KL}(P_0, P_1) \leq \frac{\log 2}{n}$ – we will choose $\delta = \sigma \sqrt{\frac{2 \log 2}{n}}$. Then, by Corollary 1, we have the following:

$$\begin{aligned} R_n^* &\geq \frac{1}{8} \Phi\left(\frac{\delta}{2}\right) \\ &= \frac{1}{8} \left(\frac{\sigma}{2} \sqrt{\frac{2 \log 2}{n}} \right)^2 \\ &= \frac{\log 2}{16} \frac{\sigma^2}{n} \end{aligned}$$

Previously, we have seen that the sample mean estimator achieves the σ^2/n rate.

2 Estimating Mean of Bernoulli Distribution

Before we proceed, we will state some facts about Bernoulli distributions that we will be using for deriving the lower bounds:

- If $P \sim \text{Bern}(p)$ and $Q \sim \text{Bern}(q)$, then we have,

$$2(p - q)^2 \underbrace{\leq}_{\textcircled{1}} \text{KL}(P, Q) \underbrace{\leq}_{\textcircled{2}} \frac{(p - q)^2}{q(1 - q)} \quad (2)$$

The proof of relations in Equation (2) above is left for exercise. The first inequality (tagged as $\textcircled{1}$) is by Pinsker's inequality, and the second inequality (tagged as $\textcircled{2}$) will use the fact that $\log(x) \leq x - 1$.

- Therefore, if q is bounded away from 0 and 1, then,

$$\text{KL}(P, Q) \in \Theta((p - q)^2)$$

This is also a general statement about sub-Gaussian distributions (i.e., $\text{KL}(P, Q) \in \Theta((\mu_P - \mu_Q)^2)$).

With this setup, let's begin our example of mean estimation for Bernoulli distributions. Let $S = \{X_1, \dots, X_n\}$ with n i.i.d samples drawn from $P \in \mathcal{P}$, where $\mathcal{P} = \{\text{Bern}(\mu \mid \mu \in [0, 1])\}$. Let $P_0 = \text{Bern}(\frac{1}{2})$ and $P_1 = \text{Bern}(\frac{1}{2} + \delta)$, then we have $|\theta(P_0) - \theta(P_1)| = \delta$. Then, by the facts introduced above, we have

$$\text{KL}(P_0, P_1) \leq \frac{(p_0 - p_1)^2}{p_1(1 - p_1)} = \frac{\delta^2}{1/4}$$

In order to satisfy the requirement of Corollary 1 of $\text{KL}(P_0, P_1) \leq \frac{\log 2}{n}$, we will choose $\delta = \frac{1}{2} \sqrt{\frac{\log 2}{n}}$. With this chosen δ , by Corollary 1, we get

$$\begin{aligned} R_n^* &\geq \frac{1}{8} \Phi\left(\frac{\delta}{2}\right) \\ &= \frac{1}{8} \left(\sqrt{\frac{\log 2}{n}} \cdot \frac{1}{4} \right)^2 \\ &= \frac{\log 2}{128} \frac{1}{n} \end{aligned}$$

Previously, we have seen that the sample mean estimator achieves $\frac{1}{n}$ rate.

3 A Simplified Regression Problem

Let $S = \{X_1, \dots, X_n\}$ where $(X_i, Y_i) \stackrel{\text{i.i.d.}}{\sim} P \in \mathcal{P}$ where $X_i \stackrel{\text{i.i.d.}}{\sim} \text{unif}([0, 1])$ and $Y_i \stackrel{\text{i.i.d.}}{\sim} \sigma$ -subGaussian distribution with mean $f(X_i)$. Here, we will assume that the function, $f : [0, 1] \rightarrow [0, 1]$, is L -Lipschitz.

Thus, by our assumption, we have

$$\mathcal{P} = \{P_{XY} \mid \underbrace{P_X = \text{unif}([0, 1])}_{\text{marginal of X}}, \underbrace{f(\cdot) = \mathbb{E}[Y|X = \cdot]}_{\text{regression function}} \text{ is } L\text{-Lipschitz with range } [0, 1], Y|X \text{ is } \sigma\text{-subgaussian}\}$$

We will begin with a simplified regression problem, where we are interested in estimating the value of the regression function at $x = 1/2$, i.e. $\theta = f(1/2) = \mathbb{E}[Y|X = 1/2]$ and the criterion for measuring ‘good’-ness is the standard squared error loss: $\Phi \circ \rho(\theta_1, \theta_2) \triangleq (\theta_1 - \theta_2)^2$.

3.1 Lower Bound

Let us first obtain the lower bound for the minimax risk. For this, we need to choose a P_0 and P_1 . To keep things simple, we will choose them to be the following:

$$P_0 : P_0(y|x) = \mathcal{N}(f_0(x), \sigma^2), \underbrace{P_0(x) = \text{unif}([0, 1])}_{\text{by assumption}}$$

$$P_1 : P_1(y|x) = \mathcal{N}(f_1(x), \sigma^2), \underbrace{P_1(x) = \text{unif}([0, 1])}_{\text{by assumption}}$$

Recall that we want the gap between two parameters to be large, i.e. large $\delta = |f_0(1/2) - f_1(1/2)|$, while ensuring that the two distributions are hard to distinguish, i.e. small $\text{KL}(P_0, P_1)$. Given that we require the regression functions, f_0 and f_1 , to be L -Lipschitz, for the simplistic setting of $f_0 \equiv 0$, the regression function f_1 would look something like the one shown below in Figure 1.




Figure 1: Here, we’d need $\delta/L \leq 1/2$ for f_1 to be well-defined.

We can mathematically define these functions, f_0 and f_1 , as follows:

$$f_0(x) \triangleq 0 \quad \forall x \in [0, 1]$$

$$f_1(x) \triangleq \begin{cases} L(x - (1/2 - \delta/L)), & \text{if } x \in [1/2 - \delta/L, 1/2] \\ L(1/2 + \delta/L - x), & \text{if } x \in [1/2, 1/2 + \delta/L] \\ 0, & \text{else} \end{cases}$$

For obtaining the lower bound as derived in Corollary 1, recall that we need to choose δ such that $KL(P_0, P_1) \leq \frac{\log 2}{n}$. To do this, let us first compute $KL(P_0, P_1)$:

$$\begin{aligned}
KL(P_0, P_1) &= \int_0^1 \int_{-\infty}^{\infty} P_0(x, y) \log \left(\frac{P_0(x, y)}{P_1(x, y)} \right) dy dx \\
&= \int_0^1 \int_{-\infty}^{\infty} P_0(y|x) \underbrace{P_0(x)}_{=1} \log \left(\frac{P_0(y|x) \overbrace{P_0(x)}^{=1}}{P_1(y|x) \underbrace{P_1(x)}_{=1}} \right) dy dx \\
&= \int_0^1 \int_{-\infty}^{\infty} P_0(y|x) \log \left(\frac{P_0(y|x)}{P_1(y|x)} \right) dy dx \\
&= \int_0^1 KL(\mathcal{N}(0, \sigma^2), \mathcal{N}(f(x), \sigma^2)) dx \\
&= \int_0^1 \frac{1}{2\sigma^2} (0 - f(x))^2 dx \quad \left(\because KL(\mathcal{N}(\mu_1, \sigma^2), \mathcal{N}(\mu_2, \sigma^2)) = \frac{1}{2\sigma^2} (\mu_1 - \mu_2)^2 \right) \\
&= \frac{1}{2\sigma^2} \left(\int_{1/2-\delta/L}^{1/2} L^2 (x - 1/2 + \delta/L)^2 dx + \int_{1/2}^{1/2+\delta/L} L^2 (1/2 + \delta/L - x)^2 dx \right) \\
&= \frac{\delta^3}{3\sigma^2 L}
\end{aligned}$$

So, we can choose $\delta = \frac{(3\sigma^2 \cdot L \cdot \log 2)^{1/3}}{n^{1/3}}$ to ensure $KL(P_0, P_1) \leq \frac{\log 2}{n}$. Application of Corollary 1 now tells us that

$$R_n^* \geq \frac{1}{8} \Phi\left(\frac{\delta}{2}\right) = \frac{1}{8} \cdot \frac{\delta^2}{4} = C \cdot \frac{\sigma^{4/3} L^{2/3}}{n^{2/3}} \quad (3)$$

where $C = \frac{1}{32} (3 \log 2)^{1/3}$.

Remark 3.1. One minor detail here is that, for our construction, since we need $\delta/L \leq 1/2$ (see Figure 1), the lower bound in Equation (3) applies only if $n \geq \frac{3\sigma^2 (\log 2)^2}{L^2}$.

Remark 3.2. Previously, when we were estimating the mean of a Gaussian distribution, we were getting a $\frac{1}{n}$ ‘rate’ in the lower bound as opposed to $\frac{1}{n^{2/3}}$ that we have obtained here for regression in Equation (3). This is because we observe data samples around the point $X = \frac{1}{2}$. We will note later that we will get similar, weaker rates even when we look at more realistic regression settings wherein we are not doing point estimation like we are here in our toy setting.

3.2 Upper Bound

Having derived the lower bound, let us derive the upper bound now. While our focus is on lower bounds, this exercise will be useful when we derive upper bounds for general regression problems later on. For this, we will rely on the idea of estimating $\theta = f(1/2)$ by calculating the average of points in a neighborhood around the point $\frac{1}{2}$. Towards this, we’ll now define a few quantities as follows:

$$\begin{aligned}
N(S) &\triangleq \sum_{i=1}^n \mathbb{I}\{X_i \in (1/2 - \alpha, 1/2 + \alpha)\} \\
\hat{\theta}(S) &\triangleq \begin{cases} 0, & \text{if } N(S) = 0 \\ \frac{1}{N(S)} \sum_{i=1}^n \mathbb{I}\{X_i \in (1/2 - \alpha, 1/2 + \alpha)\} y_i, & \text{if } N(S) > 0 \end{cases}
\end{aligned}$$

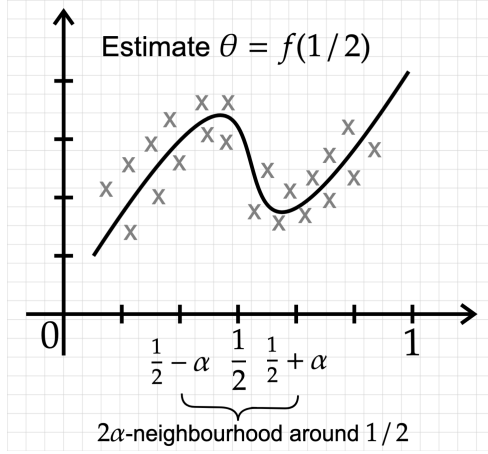


Figure 2: Idea: Estimate $\theta = f(1/2)$ by calculating the average of points in a neighborhood around the point $1/2$

Having defined the estimator, $\hat{\theta}(S)$, we now wish to compute the risk of this estimator. For this, we will define a ‘Good Event’: $G \triangleq \{N(S) \geq n\alpha\}$.

Remark 3.3 (Rationale for defining G the way we have). *If we observe points sampled uniformly between 0 and 1, the probability of each of them falling in one of the “ 2α intervals” is 2α . Thus, we see that $N(S) \sim \text{Binomial}(n, 2\alpha)$ Since, $\mathbb{E}[N(S)] = 2n\alpha$, the ‘Good Event’ is saying that ‘at least half of these points should fall near $\theta = \frac{1}{2}$ ’.*

Given the way we have defined G , it immediately leads us to the following inequality:

$$\begin{aligned} \mathbb{P}(G^c) &= \mathbb{P}\left[\sum_{i=1}^n \mathbb{I}\{X_i \in (1/2 - \alpha, 1/2 + \alpha)\} \leq n\alpha\right] \\ &\leq \exp(-2n\alpha^2) \quad (\because \text{By Hoeffding's inequality}) \end{aligned}$$

Thus, by the tower property of expectation operator, we can write the following:

$$\mathbb{E}[(\hat{\theta}(S) - \theta)^2] = \mathbb{E}[(\hat{\theta}(S) - \theta)^2 | G] \underbrace{\mathbb{P}(G)}_{\leq 1} + \underbrace{\mathbb{E}[(\hat{\theta}(S) - \theta)^2 | G^c]}_{\leq \theta^2 \leq 1} \underbrace{\mathbb{P}(G^c)}_{\leq \exp(-2n\alpha^2)} \quad (4)$$

To upper bound $\mathbb{E}[(\hat{\theta}(S) - \theta)^2 | G]$, we will expand $(\hat{\theta}(S) - \theta)^2$ as follows:

$$\begin{aligned} (\hat{\theta}(S) - \theta)^2 &= \left(\frac{1}{N(S)} \sum_{i=1}^n A_i Y_i - \theta\right)^2 = \left(\frac{1}{N(S)} \sum_{i=1}^n A_i (Y_i - \theta)\right)^2 \\ &\quad (\text{where } A_i = \mathbb{I}\{X_i \in (1/2 - \alpha, 1/2 + \alpha)\}) \\ &= \underbrace{\left(\frac{1}{N(S)} \sum_{i=1}^n A_i (Y_i - f(X_i))\right)^2}_v + \underbrace{\left(\frac{1}{N(S)} \sum_{i=1}^n A_i (f(X_i) - \theta)\right)^2}_b \end{aligned}$$

Here, we can think of the quantities v, b as the variance and bias respectively. Therefore we can write the following

$$\therefore \mathbb{E}[(\hat{\theta}(S) - \theta)^2 | G] = \mathbb{E}[(v + b)^2 | G] = \mathbb{E}[b^2 | G] + \mathbb{E}[v^2 | G] + 2\mathbb{E}[bv | G] \quad (5)$$

Calculating and bounding each of the conditional expectations separately yields the following:

$$\begin{aligned}
\mathbb{E}[v^2|G] &= \mathbb{E}\left[\mathbb{E}\left[\left(\frac{1}{N} \sum_{i=1}^{N(S)} A_i(Y_i - f(X_i))\right)^2 \middle| G, X_1, \dots, X_n\right]\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[\frac{1}{N(S)^2} \sum_{i=1}^N A_i(Y_i - f(X_i))^2 \middle| G, X_1, \dots, X_n\right]\right] \\
&\leq \mathbb{E}\left[\frac{1}{N(S)^2} \cdot N(S)\sigma^2 \middle| G\right] (\because Y_i - f(x_i) \text{ is } \sigma\text{-subGaussian}) \\
&= \mathbb{E}\left[\frac{1}{N(S)} \cdot \sigma^2 \middle| G\right] \tag{6}
\end{aligned}$$

$$\leq \frac{\sigma^2}{n\alpha} (\because N(S) \geq n\alpha \text{ under the good event } G) \tag{7}$$

Similarly, we obtain the following

$$\begin{aligned}
\mathbb{E}[b^2|G] &= \mathbb{E}\left[\left(\frac{1}{N(S)} \sum_{i=1}^{N(S)} A_i(\underbrace{f(x_i) - \theta}_{f(x_i) - f(1/2) \leq L\alpha})\right)^2 \middle| G\right] \\
&\leq L^2\alpha^2 \tag{8}
\end{aligned}$$

and

$$\begin{aligned}
\mathbb{E}[bv|G] &= \mathbb{E}\left[\left[\left(\frac{1}{n} \sum_{i=1}^n A_i(Y_i - f(X_i))\right)\left(\frac{1}{n} \sum_{i=1}^n A_i(f(X_i) - \theta)\right)\right] \middle| G, X_1, \dots, X_n\right] \\
&= 0 (\because Y_i - f(X_i)|X_i = x_i \text{ have zero mean and are independent}) \tag{9}
\end{aligned}$$

Combining Equations (4) – (9) together, we get

$$\mathbb{E}[(\hat{\theta}(S) - \theta)^2] \leq \frac{\sigma^2}{n\alpha} + L^2\alpha^2 + e^{-2n\alpha^2} \tag{10}$$

For a tighter upper bound, optimizing the R.H.S. above yields

$$\alpha^* = \frac{\sigma^{2/3}}{L^{2/3}n^{1/3}} \tag{11}$$

$$\Rightarrow \mathbb{E}[(\hat{\theta}(S) - \theta)^2] \leq \frac{2\sigma^{4/3}L^{2/3}}{n^{2/3}} + e^{-\frac{2\sigma^{4/3}}{L^{4/3}}n^{1/3}} \tag{12}$$

Remark 3.4. Note that the first term in R.H.S. of Equation (12) above is precisely the lower bound we had got earlier. But this requires that we choose α^* as per Equation (11), which requires knowledge of both σ and L , which may or may not be known.

However, if one wanted to choose α without the knowledge of σ and L , say, $\alpha = \frac{1}{N^{1/3}}$ instead of $\alpha = \alpha^*$, then we get the following loose upper bound for Equation (10) which is tight with respect to n but not σ or L :

$$\mathbb{E}[(\hat{\theta}(S) - \theta)^2] \leq e^{-2n^{1/3}} + \frac{1}{n^{2/3}}(L^2 + \sigma^2) \tag{13}$$