

Chapter 36

Minimax Theory

Minimax theory provides a rigorous framework for establishing the best possible performance of a procedure under given assumptions. In this chapter we discuss several techniques for bounding the minimax risk of a statistical problem, including the Le Cam and Fano methods.

36.1 Introduction

When solving a statistical learning problem, there are often many procedures to choose from. This leads to the following question: how can we tell if one statistical learning procedure is better than another? One answer is provided by *minimax theory* which is a set of techniques for finding the minimum, worst case behavior of a procedure.

In this chapter we rely heavily on the following sources: Yu (2008), Tsybakov (2009) and van der Vaart (1998).

36.2 Definitions and Notation

Let \mathcal{P} be a set of distributions and let X_1, \dots, X_n be a sample from some distribution $P \in \mathcal{P}$. Let $\theta(P)$ be some function of P . For example, $\theta(P)$ could be the mean of P , the variance of P or the density of P . Let $\hat{\theta} = \hat{\theta}(X_1, \dots, X_n)$ denote an estimator. Given a metric d , the *minimax risk* is

$$R_n \equiv R_n(\mathcal{P}) = \inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \quad (36.1)$$

where the infimum is over all estimators. The *sample complexity* is

$$n(\epsilon, \mathcal{P}) = \min \left\{ n : R_n(\mathcal{P}) \leq \epsilon \right\}. \quad (36.2)$$

36.3 Example. Suppose that $\mathcal{P} = \{N(\theta, 1) : \theta \in \mathbb{R}\}$ where $N(\theta, 1)$ denotes a Gaussian with mean θ and variance 1. Consider estimating θ with the metric $d(a, b) = (a - b)^2$. The minimax risk is

$$R_n = \inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[(\hat{\theta} - \theta)^2]. \quad (36.4)$$

In this example, θ is a scalar. \square

36.5 Example. Let $(X_1, Y_1), \dots, (X_n, Y_n)$ be a sample from a distribution P . Let $m(x) = \mathbb{E}_P(Y | X = x) = \int y dP(y | X = x)$ be the regression function. In this case, we might use the metric $d(m_1, m_2) = \int (m_1(x) - m_2(x))^2 dx$ in which case the minimax risk is

$$R_n = \inf_{\hat{m}} \sup_{P \in \mathcal{P}} \mathbb{E}_P \left[\int (\hat{m}(x) - m(x))^2 dx \right]. \quad (36.6)$$

In this example, θ is a function. \square

Notation. Recall that the Kullback-Leibler distance between two distributions P_0 and P_1 with densities p_0 and p_1 is defined to be

$$\text{KL}(P_0, P_1) = \int \log \left(\frac{dP_0}{dP_1} \right) dP_0 = \int \log \left(\frac{p_0(x)}{p_1(x)} \right) p_0(x) dx.$$

The appendix defines several other distances between probability distributions and explains how these distances are related. We write $a \wedge b = \min\{a, b\}$ and $a \vee b = \max\{a, b\}$. If P is a distribution with density p , the product distribution for n iid observations is P^n with density $p^n(x) = \prod_{i=1}^n p(x_i)$. It is easy to check that $\text{KL}(P_0^n, P_1^n) = n \text{KL}(P_0, P_1)$. For positive sequences a_n and b_n we write $a_n \asymp b_n$ to mean that there exists $C > 0$ such that $a_n \geq C b_n$ for all large n . $a_n \asymp b_n$ if a_n/b_n is strictly bounded away from zero and infinity for all large n ; that is, $a_n = O(b_n)$ and $b_n = O(a_n)$.

36.3 Bounding the Minimax Risk

The way we find R_n is to find an upper bound and a lower bound. To find an upper bound, let $\hat{\theta}$ be any estimator. Then

$$R_n = \inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \leq \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \equiv U_n. \quad (36.7)$$

So the maximum risk of any estimator provides an upper bound U_n . Finding a lower bound L_n is harder. We will consider two methods: the Le Cam method and the Fano method. If the lower and upper bound are close, then we have succeeded. For example, if $L_n = c n^{-\alpha}$ and $U_n = C n^{-\alpha}$ for some positive constants c, C and α , then we have established that the minimax rate of convergence is $n^{-\alpha}$.

36.4 Lower Bound Method 1: Le Cam

36.8 Theorem. Let \mathcal{P} be a set of distributions. For any pair $P_0, P_1 \in \mathcal{P}$,

$$\inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \geq \frac{\Delta}{4} \int [p_0^n(x) \wedge p_1^n(x)] dx \geq \frac{\Delta}{8} e^{-n\text{KL}(P_0, P_1)} \quad (36.9)$$

where $\Delta = d(\theta(P_0), \theta(P_1))$.

Remark: The second inequality is useful, if $\text{KL}(P_0, P_1) = \infty$, since it is usually difficult to compute $\int [p_0^n(x) \wedge p_1^n(x)] dx$ directly. An alternative is

$$\int [p_0^n(x) \wedge p_1^n(x)] dx \geq \frac{1}{2} \left(1 - \frac{1}{2} \int |p_0 - p_1| \right)^{2n}. \quad (36.10)$$

36.11 Corollary. Suppose there exist $P_0, P_1 \in \mathcal{P}$ such that $\text{KL}(P_0, P_1) \leq \log 2/n$. Then

$$\inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \geq \frac{\Delta}{16} \quad (36.12)$$

where $\Delta = d(\theta(P_0), \theta(P_1))$.

Proof. Let $\theta_0 = \theta(P_0)$, $\theta_1 = \theta(P_1)$ and $\Delta = d(\theta_0, \theta_1)$. First suppose that $n = 1$ so that we have a single observation X . An estimator $\hat{\theta}$ defines a test statistic ψ , namely,

$$\psi(X) = \begin{cases} 1 & \text{if } d(\hat{\theta}, \theta_1) \leq d(\hat{\theta}, \theta_0) \\ 0 & \text{if } d(\hat{\theta}, \theta_1) > d(\hat{\theta}, \theta_0). \end{cases}$$

If $P = P_0$ and $\psi = 1$ then

$$\Delta = d(\theta_0, \theta_1) \leq d(\theta_0, \hat{\theta}) + d(\theta_1, \hat{\theta}) \leq d(\theta_0, \hat{\theta}) + d(\theta_0, \hat{\theta}) = 2d(\theta_0, \hat{\theta})$$

and so $d(\theta_0, \hat{\theta}) \geq \frac{\Delta}{2}$. Hence

$$\mathbb{E}_{P_0}[d(\hat{\theta}, \theta_0)] \geq \mathbb{E}_{P_0}[d(\hat{\theta}, \theta_0)I(\psi = 1)] \geq \frac{\Delta}{2} \mathbb{E}_{P_0}[I(\psi = 1)] = \frac{\Delta}{2} P_0(\psi = 1). \quad (36.13)$$

Similarly,

$$\mathbb{E}_{P_1}[d(\hat{\theta}, \theta_1)] \geq \frac{\Delta}{2} P_1(\psi = 0). \quad (36.14)$$

Taking the maximum of (36.13) and (36.14), we have

$$\sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \geq \max_{P \in \{P_0, P_1\}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \geq \frac{\Delta}{2} \max\{P_0(\psi = 1), P_1(\psi = 0)\}.$$

Taking the infimum over all estimators, we have

$$\inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \geq \frac{\Delta}{2} \pi$$

where

$$\pi = \inf_{\psi} \max_{j=0,1} P_j(\psi \neq j). \quad (36.15)$$

Since a maximum is larger than an average,

$$\pi = \inf_{\psi} \max_{j=0,1} P_j(\psi \neq j) \geq \inf_{\psi} \frac{P_0(\psi \neq 0) + P_1(\psi \neq 1)}{2}.$$

The sum of the errors $P_0(\psi \neq 0) + P_1(\psi \neq 1)$ is minimized (see Lemma 36.16) by the Neyman-Pearson test

$$\psi_*(x) = \begin{cases} 0 & \text{if } p_0(x) \geq p_1(x) \\ 1 & \text{if } p_0(x) < p_1(x). \end{cases}$$

From Lemma 36.17,

$$\frac{P_0(\psi_* \neq 0) + P_1(\psi_* \neq 1)}{2} = \frac{1}{2} \int [p_0(x) \wedge p_1(x)] dx.$$

Thus we have shown that

$$\inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \geq \frac{\Delta}{4} \int [p_0(x) \wedge p_1(x)] dx.$$

Now suppose we have n observations. Then, replacing p_0 and p_1 with $p_0^n(x) = \prod_{i=1}^n p_0(x_i)$ and $p_1^n(x) = \prod_{i=1}^n p_1(x_i)$, we have

$$\inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \geq \frac{\Delta}{4} \int [p_0^n(x) \wedge p_1^n(x)] dx.$$

In Lemma 36.18 below, we show that $\int p \wedge q \geq \frac{1}{2} e^{-\text{KL}(P, Q)}$. Since $\text{KL}(P_0^n, P_1^n) = n\text{KL}(P_0, P_1)$, we have

$$\inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \geq \frac{\Delta}{8} e^{-n\text{KL}(P_0, P_1)}.$$

The result follows. \square

36.16 Lemma. *Given P_0 and P_1 , the sum of errors $P_0(\psi = 1) + P_1(\psi = 0)$ is minimized over all tests ψ by the Neyman-Pearson test*

$$\psi_*(x) = \begin{cases} 0 & \text{if } p_0(x) \geq p_1(x) \\ 1 & \text{if } p_0(x) < p_1(x). \end{cases}$$

Proof. See exercise 2. \square

36.17 Lemma. For the Neyman-Pearson test ψ_* ,

$$\frac{P_0(\psi_* \neq 0) + P_1(\psi_* \neq 1)}{2} = \frac{1}{2} \int [p_0(x) \wedge p_1(x)] dx$$

Proof. See exercise 3. \square

36.18 Lemma. For any P and Q , $\int p \wedge q \geq \frac{1}{2} e^{-\text{KL}(P,Q)}$.

Proof. First note that $\int \max(p, q) + \int \min(p, q) = 2$. Hence

$$\begin{aligned} 2 \int p \wedge q &\geq \left[2 - \int p \wedge q \right] \int p \wedge q = \int p \wedge q \int p \vee q \\ &\geq \left(\int \sqrt{(p \wedge q)(p \vee q)} \right)^2 \geq \left(\int \sqrt{pq} \right)^2 = \exp \left(2 \log \int \sqrt{pq} \right) \\ &= \exp \left(2 \log \int p \sqrt{q/p} \right) \geq \exp \left(2 \log \int p \log \sqrt{\frac{q}{p}} \right) = e^{-\text{KL}(P,Q)} \end{aligned}$$

where we used Jensen's inequality in the last inequality. \square

36.19 Example. Consider data $(X_1, Y_1), \dots, (X_n, Y_n)$ where $X_i \sim \text{Uniform}(0, 1)$, $Y_i = m(X_i) + \epsilon_i$ and $\epsilon_i \sim N(0, 1)$. Assume that

$$m \in \mathcal{M} = \left\{ m : |m(y) - m(x)| \leq L|x - y|, \text{ for all } x, y \in [0, 1] \right\}.$$

So \mathcal{P} is the set of distributions of the form $p(x, y) = p(x)p(y|x) = \phi(y - m(x))$ where $m \in \mathcal{M}$.

How well can we estimate $m(x)$ at some point x ? Without loss of generality, let's take $x = 0$ so the parameter of interest is $\theta = m(0)$. Let $d(\theta_0, \theta_1) = |\theta_0 - \theta_1|$. Let $m_0(x) = 0$ for all x . Let $0 \leq \epsilon \leq 1$ and define

$$m_1(x) = \begin{cases} L(\epsilon - x) & 0 \leq x \leq \epsilon \\ 0 & x \geq \epsilon. \end{cases}$$

Then $m_0, m_1 \in \mathcal{M}$ and $\Delta = |m_1(0) - m_0(0)| = L\epsilon$. The KL distance is

$$\begin{aligned} \text{KL}(P_0, P_1) &= \int_0^1 \int p_0(x, y) \log \left(\frac{p_0(x, y)}{p_1(x, y)} \right) dy dx \\ &= \int_0^1 \int p_0(x) p_0(y | x) \log \left(\frac{p_0(x) p_0(y | x)}{p_1(x) p_1(y | x)} \right) dy dx \\ &= \int_0^1 \int \phi(y) \log \left(\frac{\phi(y)}{\phi(y - m_1(x))} \right) dy dx \\ &= \int_0^\epsilon \int \phi(y) \log \left(\frac{\phi(y)}{\phi(y - m_1(x))} \right) dy dx \\ &= \int_0^\epsilon \text{KL}(N(0, 1), N(m_1(x), 1)) dx. \end{aligned}$$

In general, $\text{KL}(N(\mu_1, 1), N(\mu_2, 1)) = (\mu_1 - \mu_2)^2/2$. So

$$\text{KL}(P_0, P_1) = \frac{L^2}{2} \int_0^\epsilon (\epsilon - x)^2 dx = \frac{L^2 \epsilon^3}{6}.$$

Let $\epsilon = (6 \log 2 / (L^2 n))^{1/3}$. Then, $\text{KL}(P_0, P_1) = \log 2/n$ and hence, by Corollary 36.11,

$$\inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \geq \frac{\Delta}{16} = \frac{L\epsilon}{16} = \frac{L}{16} \left(\frac{6 \log 2}{L^2 n} \right)^{1/3} = \left(\frac{c}{n} \right)^{1/3}. \quad (36.20)$$

It is easy to show that the regressogram (regression histogram) $\hat{\theta} = \hat{m}(0)$ has risk

$$\mathbb{E}_P[d(\hat{\theta}, \theta(P))] \leq \left(\frac{C}{n} \right)^{1/3}.$$

Thus we have proved that

$$\inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \asymp n^{-\frac{1}{3}}. \quad (36.21)$$

The same calculations in d dimensions yield

$$\inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{\theta}, \theta(P))] \asymp n^{-\frac{1}{d+2}}. \quad (36.22)$$

On the squared scale we have

$$\inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d^2(\hat{\theta}, \theta(P))] \asymp n^{-\frac{2}{d+2}}. \quad (36.23)$$

Similar rates hold in density estimation. \square

36.5 Lower Bound Method II: Fano

For metrics like $d(f, g) = \int (f - g)^2$, Le Cam's method will usually not give a tight bound. Instead, we use Fano's method. Instead of choosing two distributions P_0, P_1 , we choose a finite set of distributions $P_1, \dots, P_N \in \mathcal{P}$.

36.24 Theorem. *Let $F = \{P_1, \dots, P_N\} \subset \mathcal{P}$. Let $\theta(P)$ be a parameter taking values in a metric space with metric d . Then*

$$\inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P \left(d(\hat{\theta}, \theta(P)) \right) \geq \frac{\alpha}{2} \left(1 - \frac{n\beta + \log 2}{\log N} \right) \quad (36.25)$$

where

$$\alpha = \min_{j \neq k} d(\theta(P_j), \theta(P_k)), \quad (36.26)$$

and

$$\beta = \max_{j \neq k} \text{KL}(P_j, P_k). \quad (36.27)$$

36.28 Corollary (Fano Minimax Bound). *Suppose there exists $F = \{P_1, \dots, P_N\} \subset \mathcal{P}$ such that $N \geq 16$ and*

$$\beta = \max_{j \neq k} \text{KL}(P_j, P_k) \leq \frac{\log N}{4n}. \quad (36.29)$$

Then

$$\inf_{\hat{\theta}} \max_{P \in \mathcal{P}} \mathbb{E}_P \left[d(\hat{\theta}, \theta(P)) \right] \geq \frac{\alpha}{4}. \quad (36.30)$$

Proof. Let $\hat{\theta}$ be any estimator and let $Z = \operatorname{argmin}_{j \in \{1, \dots, N\}} d(\hat{\theta}, \theta_j)$. For any $j \neq Z$, $d(\hat{\theta}_n(X), \theta_j) \geq \alpha/2$. Hence,

$$\mathbb{E}_{P_j} \left(d(\hat{\theta}_n, \theta(P_j)) \right) \geq \left(\frac{\alpha}{2} \right) P_j(Z \neq j)$$

and so

$$\max_j \mathbb{E}_{P_j} d(\hat{\theta}_n, \theta(P_j)) \geq \frac{\alpha}{2} \max_j P_j(Z \neq j) \geq \frac{\alpha}{2} \frac{1}{N} \sum_{j=1}^N P_j(Z \neq j).$$

By Fano's lemma (Lemma 36.83 in the appendix)

$$\frac{1}{N} \sum_{j=1}^N P_j(Z \neq j) \geq \left(1 - \frac{n\beta + \log 2}{\log N} \right).$$

Thus,

$$\inf_{\hat{\theta}} \sup_{P \in \mathcal{P}} \mathbb{E}_P \left(d(\hat{\theta}, \theta(P)) \right) \geq \inf_{\hat{\theta}} \max_{P \in \mathcal{F}} \mathbb{E}_P \left(d(\hat{\theta}, \theta(P)) \right) \geq \frac{\alpha}{2} \left(1 - \frac{n\beta + \log 2}{\log N} \right). \quad (36.31)$$

□

Hypercubes. To use Fano's method, we need to construct a finite class of distributions \mathcal{F} . Sometimes we use a set of the form

$$\mathcal{F} = \left\{ P_\omega : \omega \in \mathcal{X} \right\}$$

where

$$\mathcal{X} = \left\{ \omega = (\omega_1, \dots, \omega_m) : \omega_i \in \{0, 1\}, i = 1, \dots, m \right\}$$

which is called a hypercube. There are $N = 2^m$ distributions in \mathcal{F} . For $\omega, \nu \in \mathcal{X}$, define the *Hamming distance* $H(\omega, \nu) = \sum_{j=1}^m I(\omega_j \neq \nu_j)$.

One problem with a hypercube is that some pairs $P, Q \in \mathcal{F}$ might be very close together which will make $\alpha = \min_{j \neq k} d(\theta(P_j), \theta(P_k))$ small. This will result in a poor lower bound. We can fix this problem by pruning the hypercube. That is, we can find a subset $\mathcal{X}' \subset \mathcal{X}$ which has nearly the same number of elements as \mathcal{X} but such that each pair $P, Q \in \mathcal{F}' = \{P_\omega : \omega \in \mathcal{X}'\}$ is far apart. We call \mathcal{X}' a *pruned hypercube*. The technique for constructing \mathcal{X}' is the Varshamov-Gilbert lemma.

36.32 Lemma. (*Varshamov-Gilbert*) Let $\mathcal{X} = \{\omega = (\omega_1, \dots, \omega_N) : \omega_j \in \{0, 1\}\}$. Suppose that $N \geq 8$. There exists $\omega^0, \omega^1, \dots, \omega^M \in \mathcal{X}$ such that (i) $\omega^0 = (0, \dots, 0)$, (ii) $M \geq 2^{N/8}$ and (iii) $H(\omega^{(j)}, \omega^{(k)}) \geq N/8$ for $0 \leq j < k \leq M$. We call $\mathcal{X}' = \{\omega^0, \omega^1, \dots, \omega^M\}$ a *pruned hypercube*.

Proof. Let $D = \lfloor N/8 \rfloor$. Set $\omega^0 = (0, \dots, 0)$. Define $\mathcal{X}_0 = \mathcal{X}$ and $\mathcal{X}_1 = \{\omega \in \mathcal{X} : H(\omega, \omega^0) > D\}$. Let ω^1 be any element in \mathcal{X}_1 . Thus we have eliminated $\{\omega \in \mathcal{X} : H(\omega, \omega^0) \leq D\}$. Continue this way recursively and at the j^{th} step define $\mathcal{X}_j = \{\omega \in \mathcal{X}_{j-1} : H(\omega, \omega^{j-1}) > D\}$ where $j = 1, \dots, M$. Let n_j be the number of elements eliminated at step j , that is, the number of elements in $A_j = \{\omega \in \mathcal{X}_j : H(\omega, \omega^{(j)}) \leq D\}$. It follows that

$$n_j \leq \sum_{i=0}^D \binom{N}{i}.$$

The sets A_0, \dots, A_M partition \mathcal{X} and so $n_0 + n_1 + \dots + n_M = 2^N$. Thus,

$$(M+1) \sum_{i=0}^D \binom{N}{i} \geq 2^N.$$

Thus

$$M + 1 \geq \frac{1}{\sum_{i=0}^D 2^{-N} \binom{N}{i}} = \frac{1}{\mathbb{P}\left(\sum_{i=1}^N Z_i \leq \lfloor m/8 \rfloor\right)}$$

where Z_1, \dots, Z_N are iid Bernoulli (1/2) random variables. By Hoeffding's inequality,

$$\mathbb{P}\left(\sum_{i=1}^N Z_i \leq \lfloor m/8 \rfloor\right) \leq e^{-9N/32} < 2^{-N/4}.$$

Therefore, $M \geq 2^{N/8}$ as long as $N \geq 8$. Finally, note that, by construction, $H(\omega^j, \omega^k) \geq D + 1 \geq N/8$. \square

36.33 Example. Consider data $(X_1, Y_1), \dots, (X_n, Y_n)$ where $X_i \sim \text{Uniform}(0, 1)$, $Y_i = f(X_i) + \epsilon_i$ and $\epsilon_i \sim N(0, 1)$. (The assumption that X is uniform is not crucial.) Assume that f is in the Holder class \mathcal{F} defined by

$$\mathcal{F} = \left\{ f : |f^{(\ell)}(y) - f^{(\ell)}(x)| \leq L|x - y|^{\beta - \ell}, \text{ for all } x, y \in [0, 1] \right\}$$

where $\ell = \lfloor \beta \rfloor$. \mathcal{P} is the set of distributions of the form $p(x, y) = p(x)p(y|x) = \phi(y - m(x))$ where $f \in \mathcal{F}$. Let \mathcal{X}' be a pruned hypercube and let

$$\mathcal{F}' = \left\{ f_\omega(x) = \sum_{j=1}^m \omega_j \phi_j(x) : \omega \in \mathcal{X}' \right\}$$

where $m = \lceil cn^{\frac{1}{2\beta+1}} \rceil$, $\phi_j(x) = Lh^\beta K((x - X_j)/h)$, and $h = 1/m$. Here, K is any sufficiently smooth function supported on $(-1/2, 1/2)$. Let $d^2(f, g) = \int (f - g)^2$. Some calculations show that, for $\omega, \nu \in \mathcal{X}'$,

$$d(f_\omega, f_\nu) = \sqrt{H(\omega, \nu)} Lh^{\beta + \frac{1}{2}} \int K^2 \geq \sqrt{\frac{m}{8}} Lh^{\beta + \frac{1}{2}} \int K^2 \geq c_1 h^\beta.$$

We used the Varshamov-Gilbert result which implies that $H(\omega, \nu) \geq m/8$. Furthermore,

$$\text{KL}(P_\omega, P_\nu) \leq c_2 h^{2\beta}.$$

To apply Corollary 36.28, we need

$$\text{KL}(P_\omega, P_\nu) \leq \frac{\log N}{4n} = \frac{\log 2^{m/8}}{4n} = \frac{m}{32n} = \frac{1}{32nh}.$$

This holds if we set $h = (c/n)^{1/(2\beta+1)}$. In that case, $d(f_\omega, f_\nu) \geq c_1 h^\beta = c_1 (c/n)^{\beta/(2\beta+1)}$. Corollary 36.28 implies that

$$\inf_{\hat{f}} \sup_{P \in \mathcal{P}} \mathbb{E}_P[d(\hat{f}, f)] = \mathcal{X}(n^{-\frac{\beta}{2\beta+1}}).$$

It follows that

$$\inf_{\hat{f}} \sup_{P \in \mathcal{P}} \mathbb{E}_P \int (f - \hat{f})^2 \geq n^{-\frac{2\beta}{2\beta+1}}.$$

It can be shown that there are kernel estimators that achieve this rate of convergence. (The kernel has to be chosen carefully to take advantage of the degree of smoothness β .) A similar calculation in d dimensions shows that

$$\inf_{\hat{f}} \sup_{P \in \mathcal{P}} \mathbb{E}_P \int (f - \hat{f})^2 \geq n^{-\frac{2\beta}{2\beta+d}}.$$

□

36.6 Further Examples

36.6.1 Parametric Maximum Likelihood

For parametric models that satisfy weak regularity conditions, the maximum likelihood estimator is approximately minimax. Consider squared error loss which is squared bias plus variance. In parametric models with large samples, it can be shown that the variance term dominates the bias so the risk of the mle $\hat{\theta}$ roughly equals the variance.³⁵

$$R(\theta, \hat{\theta}) = \text{Var}_{\theta}(\hat{\theta}) + \text{bias}^2 \approx \text{Var}_{\theta}(\hat{\theta}). \quad (36.34)$$

The variance of the mle is approximately $\text{Var}(\hat{\theta}) \approx \frac{1}{nI(\theta)}$ where $I(\theta)$ is the Fisher information. Hence,

$$nR(\theta, \hat{\theta}) \approx \frac{1}{I(\theta)}. \quad (36.35)$$

For any other estimator θ' , it can be shown that for large n , $R(\theta, \theta') \geq R(\theta, \hat{\theta})$.

Here is a more precise statement, due to Hájek and Le Cam. The family of distributions $(P_{\theta} : \theta \in \Theta)$ with densities $(p_{\theta} : \theta \in \Theta)$ is *differentiable in quadratic mean* if there exists ℓ'_{θ} such that

$$\int \left(\sqrt{p_{\theta+h}} - \sqrt{p_{\theta}} - \frac{1}{2} h^T \ell'_{\theta} \sqrt{p_{\theta}} \right)^2 d\mu = o(\|h\|^2). \quad (36.36)$$

36.37 Theorem (Hájek and Le Cam). *Suppose that $(P_{\theta} : \theta \in \Theta)$ is differentiable in quadratic mean where $\Theta \subset \mathbb{R}^k$ and that the Fisher information I_{θ} is nonsingular. Let ψ be differentiable. Then $\psi(\hat{\theta}_n)$, where $\hat{\theta}_n$ is the mle, is asymptotically, locally, uniformly minimax in the sense that, for any estimator T_n , and any bowl-shaped ℓ ,*

$$\sup_{I \in \mathcal{I}} \liminf_{n \rightarrow \infty} \sup_{h \in I} \mathbb{E}_{\theta+h/\sqrt{n}} \ell \left(\sqrt{n} \left(T_n - \psi \left(\theta + \frac{h}{\sqrt{n}} \right) \right) \right) \geq \mathbb{E}(\ell(U)) \quad (36.38)$$

where \mathcal{I} is the class of all finite subsets of \mathbb{R}^k and $U \sim N(0, \psi'_{\theta} I_{\theta}^{-1} (\psi'_{\theta})^T)$.

³⁵Typically, the squared bias is order $O(n^{-2})$ while the variance is of order $O(n^{-1})$.

For a proof, see van der Vaart (1998). Note that the right hand side of the displayed formula is the risk of the mle. In summary: in well-behaved parametric models, with large samples, the mle is approximately minimax. There is a crucial caveat: these results break down when the number of parameters is large.

36.6.2 Estimating a Smooth Density

Here we use the general strategy to derive the minimax rate of convergence for estimating a smooth density. (See Yu (2008) for more details.)

Let \mathcal{F} be all probability densities f on $[0, 1]$ such that

$$0 < c_0 \leq f(x) \leq c_1 < \infty, \quad |f''(x)| \leq c_2 < \infty.$$

We observe $X_1, \dots, X_n \sim P$ where P has density $f \in \mathcal{F}$. We will use the squared Hellinger distance $d^2(f, g) = \int_0^1 (\sqrt{f(x)} - \sqrt{g(x)})^2 dx$ as a loss function.

Upper Bound. Let \hat{f}_n be the kernel estimator with bandwidth $h = n^{-1/5}$. Then, using bias-variance calculations, we have that

$$\sup_{f \in \mathcal{F}} \mathbb{E}_f \left(\int (\hat{f}_n(x) - f(x))^2 dx \right) \leq Cn^{-4/5}$$

for some C . But

$$\int (\sqrt{f(x)} - \sqrt{g(x)})^2 dx = \int \left(\frac{f(x) - g(x)}{\sqrt{f(x)} + \sqrt{g(x)}} \right)^2 dx \leq C' \int (f(x) - g(x))^2 dx \quad (36.39)$$

for some C' . Hence $\sup_f \mathbb{E}_f (d^2(f, \hat{f}_n)) \leq Cn^{-4/5}$ which gives us an upper bound.

Lower Bound. For the lower bound we use Fano's inequality. Let g be a bounded, twice differentiable function on $[-1/2, 1/2]$ such that

$$\int_{-1/2}^{1/2} g(x) dx = 0, \quad \int_{-1/2}^{1/2} g^2(x) dx = a > 0, \quad \int_{-1/2}^{1/2} (g'(x))^2 dx = b > 0.$$

Fix an integer m and for $j = 1, \dots, m$ define $x_j = (j - (1/2))/m$ and

$$g_j(x) = \frac{c}{m^2} g(m(x - x_j))$$

for $x \in [0, 1]$ where c is a small positive constant. Let \mathcal{M} denote the Varshamov-Gilbert pruned version of the set

$$\left\{ f_\tau = 1 + \sum_{j=1}^m \tau_j g_j(x) : \tau = (\tau_1, \dots, \tau_m) \in \{-1, +1\}^m \right\}.$$

For $f_\tau \in \mathcal{M}$, let f_τ^n denote the product density for n observations and let $\mathcal{M}_n = \left\{ f_\tau^n : f_\tau \in \mathcal{M} \right\}$. Some calculations show that, for all τ, τ' ,

$$\text{KL}(f_\tau^n, f_{\tau'}^n) = n\text{KL}(f_\tau, f_{\tau'}) \leq \frac{C_1 n}{m^4} \equiv \beta. \quad (36.40)$$

By Lemma 36.32, we can choose a subset F of \mathcal{M} with $N = e^{c_0 m}$ elements (where c_0 is a constant) and such that

$$d^2(f_\tau, f_{\tau'}) \geq \frac{C_2}{m^4} \equiv \alpha \quad (36.41)$$

for all pairs in F . Choosing $m = cn^{1/5}$ gives $\beta \leq \log N/4$ and $d^2(f_\tau, f_{\tau'}) \geq \frac{C_2}{n^{4/5}}$. Fano's lemma implies that

$$\max_j \mathbb{E}_j d^2(\hat{f}, f_j) \geq \frac{C}{n^{4/5}}.$$

Hence the minimax rate is $n^{-4/5}$ which is achieved by the kernel estimator. Thus we have shown that $R_n(\mathcal{P}) \asymp n^{-4/5}$.

This result can be generalized to higher dimensions and to more general measures of smoothness. Since the proof is similar to the one dimensional case, we state the result without proof.

36.42 Theorem. *Let \mathcal{Z} be a compact subset of \mathbb{R}^d . Let $\mathcal{F}(p, C)$ denote all probability density functions on \mathcal{Z} such that*

$$\int \sum \left| \frac{\partial^p}{\partial z_1^{p_1} \dots \partial z_d^{p_d}} f(z) \right|^2 dz \leq C$$

where the sum is over all p_1, \dots, p_d such that $\sum_j p_j = p$. Then there exists a constant $D > 0$ such that

$$\inf_{\hat{f}} \sup_{f \in \mathcal{F}(p, C)} \mathbb{E}_f \int (\hat{f}_n(z) - f(z))^2 dz \geq D \left(\frac{1}{n} \right)^{\frac{2p}{2p+1}}. \quad (36.43)$$

The kernel estimator (with an appropriate kernel) with bandwidth $h_n = n^{-1/(2p+d)}$ achieves this rate of convergence.

36.6.3 Minimax Classification

Let us now turn to classification. We focus on some results of Yang (1999), Tsybakov (2004), Mammen and Tsybakov (1999), Audibert and Tsybakov (2005) and Tsybakov and van de Geer (2005).

The data are $Z = (X_1, Y_1), \dots, (X_n, Y_n)$ where $Y_i \in \{0, 1\}$. Recall that a classifier is a function of the form $h(x) = I(x \in G)$ for some set G . The classification risk is

$$R(G) = \mathbb{P}(Y \neq h(X)) = \mathbb{P}(Y \neq I(X \in G)) = \mathbb{E}(Y - I(X \in G))^2. \quad (36.44)$$

The optimal classifier is $h^*(x) = I(x \in G^*)$ where $G^* = \{x : m(x) \geq 1/2\}$ and $m(x) = \mathbb{E}(Y | X = x)$. We are interested in how close $R(G)$ is to $R(G^*)$. Following Tsybakov (2004) we define

$$d(G, G^*) = R(G) - R(G^*) = 2 \int_{G \Delta G^*} \left| m(x) - \frac{1}{2} \right| dP_X(x) \quad (36.45)$$

where $A \Delta B = (A \cap B^c) \cup (A^c \cap B)$ and P_X is the marginal distribution of X .

There are two common types of classifiers. The first type are *plug-in classifiers* of the form $\hat{h}(x) = I(\hat{m}(x) \geq 1/2)$ where \hat{m} is an estimate of the regression function. The second type are *empirical risk minimizers* where \hat{h} is taken to be the h that minimizes the observed error rate $n^{-1} \sum_{i=1}^n (Y_i \neq h(X_i))$ as h varies over a set of classifiers \mathcal{H} . Sometimes one minimizes the error rate plus a penalty term.

According to Yang (1999), the classification problem has, under weak conditions, the same order of difficulty (in terms of minimax rates) as estimating the regression function $m(x)$. Therefore the rates are given in Example 36.75. According to Tsybakov (2004) and Mammen and Tsybakov (1999), classification is easier than regression. The apparent discrepancy is due to differing assumptions.

To see that classification error cannot be harder than regression, note that for any \hat{m} and corresponding \hat{G}

$$d(G, \hat{G}) = 2 \int_{G \Delta \hat{G}} \left| m(x) - \frac{1}{2} \right| dP_X(x) \quad (36.46)$$

$$\leq 2 \int |\hat{m}(x) - m(x)| dP_X(x) \leq 2 \sqrt{\int (\hat{m}(x) - m(x))^2 dP_X(x)} \quad (36.47)$$

so the rate of convergence of $d(G, G^*)$ is at least as fast as the regression function.

Instead of putting assumptions on the regression function m , Mammen and Tsybakov (1999) put an entropy assumption on the set of decision sets \mathcal{G} . They assume

$$\log N(\epsilon, \mathcal{G}, d) \leq A\epsilon^{-\rho} \quad (36.48)$$

where $N(\epsilon, \mathcal{G}, d)$ is the smallest number of balls of radius ϵ required to cover \mathcal{G} . They show that, if $0 < \rho < 1$, then there are classifiers with rate

$$\sup_P \mathbb{E}(d(\hat{G}, G^*)) = O(n^{-1/2}) \quad (36.49)$$

independent of dimension d . Moreover, if we add the margin (or low noise) assumption

$$\mathbb{P}_X \left(0 < \left| m(X) - \frac{1}{2} \right| \leq t \right) \leq Ct^\alpha \quad \text{for all } t > 0 \quad (36.50)$$

we get

$$\sup_P \mathbb{E}(d(\hat{G}, G^*)) = O \left(n^{-(1+\alpha)/(2+\alpha+\alpha\rho)} \right) \quad (36.51)$$

which can be nearly $1/n$ for large α and small ρ . The classifiers can be taken to be plug-in estimators using local polynomial regression. Moreover, they show that this rate is minimax. We will discuss classification in the low noise setting in more detail in another chapter.

36.6.4 Estimating a Large Covariance Matrix

Let X_1, \dots, X_n be iid Gaussian vectors of dimension d . Let $\Sigma = (\sigma_{ij})_{1 \leq i, j \leq d}$ be the $d \times d$ covariance matrix for X_i . Estimating Σ when d is large is very challenging. Sometimes we can take advantage of special structure. Bickel and Levina (2008) considered the class of covariance matrices Σ whose entries have polynomial decay. Specifically, $\Theta = \Theta(\alpha, \epsilon, M)$ is all covariance matrices Σ such that $0 < \epsilon \leq \lambda_{\min}(\Sigma) \leq \lambda_{\max}(\Sigma) \leq 1/\epsilon$ and such that

$$\max_j \sum_i \left\{ |\sigma_{ij}| : |i - j| > k \right\} \leq Mk^{-\alpha}$$

for all k . The loss function is $\|\widehat{\Sigma} - \Sigma\|$ where $\|\cdot\|$ is the operator norm

$$\|A\| = \sup_{x: \|x\|_2=1} \|Ax\|_2.$$

Bickel and Levina (2008) constructed an estimator that that converges at rate $(\log d/n)^{\alpha/(\alpha+1)}$. Cai, Zhang and Zhou (2009) showed that the minimax rate is

$$\min \left\{ n^{-2\alpha/(2\alpha+1)} + \frac{\log d}{n}, \frac{d}{n} \right\}$$

so the Bickel-Levina estimator is not rate minimax. Cai, Zhang and Zhou then constructed an estimator that is rate minimax.

36.6.5 Semisupervised Prediction

Suppose we have data $(X_1, Y_1), \dots, (X_n, Y_n)$ for a classification or regression problem. In addition, suppose we have extra unlabelled data X_{n+1}, \dots, X_N . Methods that make use of the unlabeled are called *semisupervised methods*. We discuss semisupervised methods in another Chapter.

When do the unlabeled data help? Two minimax analyses have been carried out to answer that question, namely, Lafferty and Wasserman (2007) and Singh, Nowak and Zhu (2008). Here we briefly summarize the results of the latter.

Suppose we want to estimate $m(x) = \mathbb{E}(Y | X = x)$ where $x \in \mathbb{R}^d$ and $y \in \mathbb{R}$. Let p be the density of X . To use the unlabelled data we need to link m and p in some way. A common assumption is the *cluster assumption*: m is smooth over clusters of the marginal $p(x)$. Suppose that p has clusters separated by a amount γ and that m is α smooth over each cluster. Singh, Nowak and Zhu (2008) obtained the following upper and lower minimax bounds as γ varies in 6 zones which we label I to VI. These zones relate the size of γ and the number of unlabeled points:

γ	semisupervised upper bound	supervised lower bound	unlabelled data help?
I	$n^{-2\alpha/(2\alpha+d)}$	$n^{-2\alpha/(2\alpha+d)}$	NO
II	$n^{-2\alpha/(2\alpha+d)}$	$n^{-2\alpha/(2\alpha+d)}$	NO
III	$n^{-2\alpha/(2\alpha+d)}$	$n^{-1/d}$	YES
IV	$n^{-1/d}$	$n^{-1/d}$	NO
V	$n^{-2\alpha/(2\alpha+d)}$	$n^{-1/d}$	YES
VI	$n^{-2\alpha/(2\alpha+d)}$	$n^{-1/d}$	YES

The important message is that there are precise conditions when the unlabeled data help and conditions when the unlabeled data do not help. These conditions arise from computing the minimax bounds.

36.6.6 Graphical Models

Elsewhere in the book, we discuss the problem of estimating graphical models. Here, we shall briefly mention some minimax results for this problem. Let X be a random vector from a multivariate Normal distribution P with mean vector μ and covariance matrix Σ . Note that X is a random vector of length d , that is, $X = (X_1, \dots, X_d)^T$. The $d \times d$ matrix $\mathcal{X} = \Sigma^{-1}$ is called the precision matrix. There is one node for each component of X . The undirected graph associated with P has no edge between X_j and X_k if and only if $\mathcal{X}_{jk} = 0$. The edge set is $E = \{(j, k) : \mathcal{X}_{jk} \neq 0\}$. The graph is $G = (V, E)$ where $V = \{1, \dots, d\}$ and E is the edge set. Given a random sample of vectors $X^1, \dots, X^n \sim P$ we want to estimate G . (Only the edge set needs to be estimated; the nodes are known.)

Wang, Wainwright and Ramchandran (2010) found the minimax risk for estimating G under zero-one loss. Let $\mathcal{G}_{d,r}(\lambda)$ denote all the multivariate Normals whose graphs have edge sets with degree at most r and such that

$$\min_{(i,j) \in E} \frac{|\mathcal{X}_{jk}|}{\sqrt{\mathcal{X}_{jj}\mathcal{X}_{kk}}} \geq \lambda.$$

The sample complexity $n(d, r, \lambda)$ is the smallest sample size n needed to recover the true graph with high probability. They show that for any $\lambda \in [0, 1/2]$,

$$n(d, r, \lambda) > \max \left\{ \frac{\log \binom{d-r}{2} - 1}{4\lambda^2}, \frac{\log \binom{d}{r} - 1}{\frac{1}{2} \left(\log \left(1 + \frac{r\lambda}{1-\lambda} \right) - \frac{r\lambda}{1+(r-1)\lambda} \right)} \right\}. \tag{36.52}$$

Thus, assuming $\lambda \approx 1/r$, we get that $n \geq Cr^2 \log(d - r)$.

36.6.7 Deconvolution and Measurement Error

A problem that seems to have received little attention in the machine learning literature is *deconvolution*. Suppose that $X_1, \dots, X_n \sim P$ where P has density p . We have seen that

the minimax rate for estimating p in squared error loss is $n^{-\frac{2\beta}{2\beta+1}}$ where β is the assumed amount of smoothness. Suppose we cannot observe X_i directly but instead we observe X_i with error. Thus, we observe Y_1, \dots, Y_n where

$$Y_i = X_i + \epsilon_i, \quad i = 1, \dots, n. \quad (36.53)$$

The minimax rates for estimating p change drastically. A good account is given in Fan (1991). As an example, if the noise ϵ_i is Gaussian, then Fan shows that the minimax risk satisfies

$$R_n \geq C \left(\frac{1}{\log n} \right)^\beta$$

which means that the problem is essentially hopeless.

Similar results hold for nonparametric regression. In the usual nonparametric regression problem we observe $Y_i = m(X_i) + \epsilon_i$ and we want to estimate the function m . If we observe $X_i^* = X_i + \delta_i$ instead of X_i then again the minimax rates change drastically and are logarithmic of the δ_i 's are Normal (Fan and Truong 1993). This is known as *measurement error* or *errors in variables*.

This is an interesting example where minimax theory reveals surprising and important insight.

36.6.8 Normal Means

Perhaps the best understood cases in minimax theory involve normal means. First suppose that $X_1, \dots, X_n \sim N(\theta, \sigma^2)$ where σ^2 is known. A function g is *bowl-shaped* if the sets $\{x : g(x) \leq c\}$ are convex and symmetric about the origin. We will say that a loss function ℓ is bowl-shaped if $\ell(\theta, \hat{\theta}) = g(\theta - \hat{\theta})$ for some bowl-shaped function g .

36.54 Theorem. *The unique³⁶ estimator that is minimax for every bowl-shaped loss function is the sample mean \bar{X}_n .*

For a proof, see Wolfowitz (1950).

Now consider estimating several normal means. Let $X_j = \theta_j + \epsilon_j/\sqrt{n}$ for $j = 1, \dots, n$ and suppose we want to estimate $\theta = (\theta_1, \dots, \theta_n)$ with loss function $\ell(\hat{\theta}, \theta) = \sum_{j=1}^n (\hat{\theta}_j - \theta_j)^2$. Here, $\epsilon_1, \dots, \epsilon_n \sim N(0, \sigma^2)$. This is called the *normal means problem*.

There are strong connections between the normal means problem and nonparametric learning. For example, suppose we want to estimate a regression function $f(x)$ and we observe data $Z_i = f(i/n) + \delta_i$ where $\delta_i \sim N(0, \sigma^2)$. Expand f in an orthonormal basis: $f(x) = \sum_j \theta_j \psi_j(x)$. An estimate of θ_j is $X_j = \frac{1}{n} \sum_{i=1}^n Z_i \psi_j(i/n)$. It follows that $X_j \approx N(\theta_j, \sigma^2/n)$. This connection can be made very rigorous; see Brown and Low (1996).

³⁶Up to sets of measure 0.

The minimax risk depends on the assumptions about θ .

36.55 Theorem (Pinsker).

1. If $\Theta_n = \mathbb{R}^n$ then $R_n = \sigma^2$ and $\hat{\theta} = X = (X_1, \dots, X_n)$ is minimax.
2. If $\Theta_n = \{\theta : \sum_j^n \theta_j^2 \leq C^2\}$ then

$$\liminf_{n \rightarrow \infty} \inf_{\hat{\theta}} \sup_{\theta \in \Theta_n} R(\hat{\theta}, \theta) = \frac{\sigma^2 C^2}{\sigma^2 + C^2}. \quad (36.56)$$

Define the James-Stein estimator

$$\hat{\theta}_{\text{JS}} = \left(1 - \frac{(n-2)\sigma^2}{\frac{1}{n} \sum_{j=1}^n X_j^2} \right) X. \quad (36.57)$$

Then

$$\lim_{n \rightarrow \infty} \sup_{\theta \in \Theta_n} R(\hat{\theta}_{\text{JS}}, \theta) = \frac{\sigma^2 C^2}{\sigma^2 + C^2}. \quad (36.58)$$

Hence, $\hat{\theta}_{\text{JS}}$ is asymptotically minimax.

3. Let $X_j = \theta_j + \epsilon_j$ for $j = 1, 2, \dots$, where $\epsilon_j \sim N(0, \sigma^2/n)$.

$$\Theta = \left\{ \theta : \sum_{j=1}^{\infty} \theta_j^2 a_j^2 \leq C^2 \right\} \quad (36.59)$$

where $a_j^2 = (\pi j)^{2p}$. Let R_n denote the minimax risk. Then

$$\min_{n \rightarrow \infty} n^{\frac{2p}{2p+1}} R_n = \left(\frac{\sigma}{\pi} \right)^{\frac{2p}{2p+1}} C^{\frac{2}{2p+1}} \left(\frac{p}{p+1} \right)^{\frac{2p}{2p+1}} (2p+1)^{\frac{1}{2p+1}}. \quad (36.60)$$

Hence, $R_n \asymp n^{-\frac{2p}{2p+1}}$. An asymptotically minimax estimator is the Pinsker estimator defined by $\hat{\theta} = (w_1 X_1, w_2 X_2, \dots)$ where $w_j = [1 - (a_j/\mu)]_+$ and μ is determined by the equation

$$\frac{\sigma^2}{n} \sum_j a_j (\mu - a_j)_+ = C^2.$$

The set Θ in (36.59) is called a *Sobolev ellipsoid*. This set corresponds to smooth functions in the function estimation problem. The Pinsker estimator corresponds to estimating a function by smoothing. The main message to take away from all of this is that minimax estimation under smoothness assumptions requires shrinking the data appropriately.

36.7 Adaptation

The results in this chapter provide minimax rates of convergence and estimators that achieve these rates. However, the estimators depend on the assumed parameter space. For example, estimating a β -times differential regression function requires using an estimator tailored to the assumed amount of smoothness to achieve the minimax rate $n^{-\frac{2\beta}{2\beta+1}}$. There are estimators that are *adaptive*, meaning that they achieve the minimax rate without the user having to know the amount of smoothness. See, for example, Chapter 9 of Wasserman (2006) and the references therein.

36.8 The Bayesian Connection

Another way to find the minimax risk and to find a minimax estimator is to use a carefully constructed Bayes estimator. In this section we assume we have a parametric family of densities $\{p(x; \theta) : \theta \in \Theta\}$ and that our goal is to estimate the parameter θ . Since the distributions are indexed by θ , we can write the risk as $R(\theta, \hat{\theta}_n)$ and the maximum risk as $\sup_{\theta \in \Theta} R(\theta, \hat{\theta}_n)$.

Let Q be a prior distribution for θ . The *Bayes risk (with respect to Q)* is defined to be

$$B_Q(\hat{\theta}_n) = \int R(\theta, \hat{\theta}_n) dQ(\theta) \quad (36.61)$$

and the *Bayes estimator with respect to Q* is the estimator $\bar{\theta}_n$ that minimizes $B_Q(\hat{\theta}_n)$. For simplicity, assume that Q has a density q . The posterior density is then

$$q(\theta | X^n) = \frac{p(X_1, \dots, X_n; \theta) q(\theta)}{m(X_1, \dots, X_n)}$$

where $m(x_1, \dots, x_n) = \int p(x_1, \dots, x_n; \theta) q(\theta) d\theta$.

36.62 Lemma. *The Bayes risk can be written as*

$$\int \left(\int L(\theta, \hat{\theta}_n) q(\theta | x_1, \dots, x_n) d\theta \right) m(x_1, \dots, x_n) dx_1 \cdots dx_n.$$

Proof. See exercise 15. \square

It follows from this lemma that the Bayes estimator can be obtained by finding $\hat{\theta}_n = \hat{\theta}_n(x_1, \dots, x_n)$ to minimize the inner integral $\int L(\theta, \hat{\theta}_n) q(\theta | x_1, \dots, x_n) d\theta$. Often, this is an easy calculation.

36.63 Example. Suppose that $L(\theta, \hat{\theta}_n) = (\theta - \hat{\theta}_n)^2$. Then the Bayes estimator is the posterior mean $\bar{\theta}_Q = \int \theta q(\theta | x_1, \dots, x_n) d\theta$. \square

Now we link Bayes estimators to minimax estimators.

36.64 Theorem. Let $\hat{\theta}_n$ be an estimator. Suppose that (i) the risk function $R(\theta, \hat{\theta}_n)$ is constant as a function of θ and (ii) $\hat{\theta}_n$ is the Bayes estimator for some prior Q . Then $\hat{\theta}_n$ is minimax.

Proof. We will prove this by contradiction. Suppose that $\hat{\theta}_n$ is not minimax. Then there is some other estimator θ' such that

$$\sup_{\theta \in \Theta} R(\theta, \theta') < \sup_{\theta \in \Theta} R(\theta, \hat{\theta}_n). \tag{36.65}$$

Now,

$$\begin{aligned} B_Q(\theta') &= \int R(\theta, \theta') dQ(\theta) && \text{definition of Bayes risk} \\ &\leq \sup_{\theta \in \Theta} R(\theta, \theta') && \text{average is less than sup} \\ &< \sup_{\theta \in \Theta} R(\theta, \hat{\theta}_n) && \text{from (36.65)} \\ &= \int R(\theta, \hat{\theta}_n) dQ(\theta) && \text{since risk is constant} \\ &= B_Q(\hat{\theta}_n) && \text{definition of Bayes risk.} \end{aligned}$$

So $B_Q(\theta') < B_Q(\hat{\theta}_n)$. This is a contradiction because $\hat{\theta}_n$ is the Bayes estimator for Q so it must minimize B_Q . \square

36.66 Example. Let $X \sim \text{Binomial}(n, \theta)$. The mle is X/n . Let $L(\theta, \hat{\theta}_n) = (\theta - \hat{\theta}_n)^2$. Define

$$\hat{\theta}_n = \frac{\frac{X}{n} + \sqrt{\frac{1}{4n}}}{1 + \sqrt{\frac{1}{n}}}.$$

Some calculations show that this is the posterior mean under a $\text{Beta}(\alpha, \beta)$ prior with $\alpha = \beta = \sqrt{n/4}$. By computing the bias and variance of $\hat{\theta}_n$ it can be seen that $R(\theta, \hat{\theta}_n)$ is constant. Since $\hat{\theta}_n$ is Bayes and has constant risk, it is minimax. \square

36.67 Example. Let us now show that the sample mean is minimax for the Normal model. Let $X \sim N_p(\theta, I)$ be multivariate Normal with mean vector $\theta = (\theta_1, \dots, \theta_p)$. We will prove that $\hat{\theta}_n = X$ is minimax when $L(\theta, \hat{\theta}_n) = \|\hat{\theta}_n - \theta\|^2$. Assign the prior $Q = N(0, c^2 I)$. Then the posterior is

$$N\left(\frac{c^2 x}{1 + c^2}, \frac{c^2}{1 + c^2} I\right). \tag{36.68}$$

The Bayes risk $B_Q(\hat{\theta}_n) = \int R(\theta, \hat{\theta}_n) dQ(\theta)$ is minimized by the posterior mean $\tilde{\theta} = c^2 X / (1 + c^2)$. Direct computation shows that $B_Q(\tilde{\theta}) = pc^2 / (1 + c^2)$. Hence, if θ^* is any estimator, then

$$\frac{pc^2}{1 + c^2} = B_Q(\tilde{\theta}) \leq B_Q(\theta^*) = \int R(\theta^*, \theta) dQ(\theta) \leq \sup_{\theta} R(\theta^*, \theta).$$

This shows that $R(\Theta) \geq pc^2 / (1 + c^2)$ for every $c > 0$ and hence

$$R(\Theta) \geq p. \quad (36.69)$$

But the risk of $\hat{\theta}_n = X$ is p . So, $\hat{\theta}_n = X$ is minimax. \square

36.9 Nonparametric Maximum Likelihood and the Le Cam Equation

In some cases, the minimax rate can be found by finding ϵ to solve the equation $H(\epsilon) = n\epsilon^2$ where $H(\epsilon) = \log N(\epsilon)$ and $N(\epsilon)$ is the smallest number of balls of size ϵ in the Hellinger metric needed to cover \mathcal{P} . $H(\epsilon)$ is called the Hellinger entropy of \mathcal{P} . The equation $H(\epsilon) = n\epsilon^2$ is known as the *Le Cam equation*. In this section we consider one case where this is true. For more general versions of this argument, see Shen and Wong (1995), Barron and Yang (1999) and Birgé and Massart (1993).

Our goal is to estimate the density function using maximum likelihood. The loss function is Hellinger distance. Let \mathcal{P} be a set of probability density functions. We have in mind the nonparametric situation where \mathcal{P} does not correspond to some finite dimensional parametric family. Let $N(\epsilon)$ denote the Hellinger covering number of \mathcal{P} . We will make the following assumptions:

(A1) We assume that there exist $0 < c_1 < c_2 < \infty$ such that $c_1 \leq p(x) \leq c_2$ for all x and all $p \in \mathcal{P}$.

(A2) We assume that there exists $a > 0$ such that

$$H(a\epsilon, \mathcal{P}, h) \leq \sup_{p \in \mathcal{P}} H(\epsilon, B(p, 4\epsilon), h)$$

where $B(p, \delta) = \{q : h(p, q) \leq \delta\}$.

(A3) We assume $\sqrt{n}\epsilon_n \rightarrow \infty$ as $n \rightarrow \infty$ where $H(\epsilon_n) \asymp n\epsilon_n^2$.

Assumption (A1) is a very strong and is made only to make the proofs simpler. Assumption (A2) says that the local entropy and global entropy are of the same order. This is typically true in nonparametric models. Assumption (A3) says that the rate of convergence

is slower than $O(1/\sqrt{n})$ which again is typical of nonparametric problems. An example of a class \mathcal{P} that satisfies these conditions is

$$\mathcal{P} = \left\{ p : [0, 1] \rightarrow [c_1, c_2] : \int_0^1 p(x) dx = 1, \int_0^1 (p''(x))^2 dx \leq C^2 \right\}.$$

Thanks to (A1) we have,

$$\begin{aligned} \text{KL}(p, q) &\leq \chi^2(p, q) = \int \frac{(p-q)^2}{p} \leq \frac{1}{c_1} \int (p-q)^2 \\ &= \frac{1}{c_1} \int (\sqrt{p} - \sqrt{q})^2 (\sqrt{p} + \sqrt{q})^2 \\ &\leq \frac{4c_2}{c_1} \int (\sqrt{p} - \sqrt{q})^2 = Ch^2(p, q) \end{aligned} \quad (36.70)$$

where $C = 4c_2/c_1$.

Let ϵ_n solve the Le Cam equation. More precisely, let

$$\epsilon_n = \min \left\{ \epsilon : H \left(\frac{\epsilon}{\sqrt{2C}} \right) \leq \frac{n\epsilon^2}{16C} \right\}. \quad (36.71)$$

We will show that ϵ_n is the minimax rate.

Upper Bound. To show the upper bound, we will find an estimator that achieves the rate. Let $\mathcal{P}_n = \{p_1, \dots, p_N\}$ be an $\epsilon_n/\sqrt{2C}$ covering set where $N = N(\epsilon_n/\sqrt{2C})$. The set \mathcal{P}_n is an approximation to \mathcal{P} that grows with sample size n . Such a set is called *sieve*. Let \hat{p} be the mle over \mathcal{P}_n , that is, $\hat{p} = \operatorname{argmax}_{p \in \mathcal{P}_n} \mathcal{L}(p)$ where $\mathcal{L}(p) = \prod_{i=1}^n p(X_i)$ is the likelihood function. We call \hat{p} , a *sieve maximum likelihood estimator*. It is crucial that the estimator is computed over \mathcal{P}_n rather than over \mathcal{P} to prevent overfitting. Using a sieve is a type of regularization. We need the following lemma.

36.72 Lemma (Wong and Shen). *Let p_0 and p be two densities and let $\delta = h(p_0, p)$. Let Z_1, \dots, Z_n be a sample from p_0 . Then*

$$\mathbb{P} \left(\frac{\mathcal{L}(p)}{\mathcal{L}(p_0)} > e^{-n\delta^2/2} \right) \leq e^{-n\delta^2/4}.$$

Proof.

$$\begin{aligned}
\mathbb{P}\left(\frac{\mathcal{L}(p)}{\mathcal{L}(p_0)} > e^{-n\delta^2/2}\right) &= \mathbb{P}\left(\prod_{i=1}^n \sqrt{\frac{p(Z_i)}{p_0(Z_i)}} > e^{-n\delta^2/4}\right) \leq e^{n\delta^2/4} \mathbb{E}\left(\prod_{i=1}^n \sqrt{\frac{p(Z_i)}{p_0(Z_i)}}\right) \\
&= e^{n\delta^2/4} \left(\mathbb{E}\left(\sqrt{\frac{p(Z_i)}{p_0(Z_i)}}\right)\right)^n = e^{n\delta^2/4} \left(\int \sqrt{p_0 p}\right)^n \\
&= e^{n\delta^2/4} \left(1 - \frac{h^2(p_0, p)}{2}\right)^n = e^{n\delta^2/4} \exp\left(n \log\left(1 - \frac{h^2(p_0, p)}{2}\right)\right) \\
&\leq e^{n\delta^2/4} e^{-nh^2(p_0, p)/2} = e^{-n\delta^2/4}.
\end{aligned}$$

□

In what follows, we use c, c_1, c_2, \dots , to denote various positive constants.

36.73 Theorem. $\sup_{P \in \mathcal{P}} \mathbb{E}_p(h(p, \hat{p})) = O(\epsilon_n)$.

Proof. Let p_0 denote the true density. Let p_* be the element of \mathcal{P}_n that minimizes $\text{KL}(p_0, p_j)$. Hence, $\text{KL}(p_0, p_*) \leq Cd^2(p_0, p_*) \leq C(\epsilon_n^2/(2C)) = \epsilon_n^2/2$. Let

$$B = \{p \in \mathcal{P}_n : d(p_*, p) > A\epsilon_n\}$$

where $A = 1/\sqrt{2C}$. Then

$$\begin{aligned}
\mathbb{P}(h(\hat{p}, p_0) > D\epsilon_n) &\leq \mathbb{P}(h(\hat{p}, p_*) + h(p_0, p_*) > D\epsilon_n) \leq \mathbb{P}(h(\hat{p}, p_*) + \frac{\epsilon}{\sqrt{2C}} > D\epsilon_n) \\
&= \mathbb{P}(h(\hat{p}, p_*) > A\epsilon_n) = \mathbb{P}(\hat{p} \in B) \leq \mathbb{P}\left(\sup_{p \in B} \frac{\mathcal{L}(p)}{\mathcal{L}(p_*)} > 1\right) \\
&\leq \mathbb{P}\left(\sup_{p \in B} \frac{\mathcal{L}(p)}{\mathcal{L}(p_*)} > e^{-n\epsilon_n^2(A^2/2+1)}\right) \\
&\leq \mathbb{P}\left(\sup_{p \in B} \frac{\mathcal{L}(p)}{\mathcal{L}(p_*)} > e^{-n\epsilon_n^2 A^2/2}\right) + \mathbb{P}\left(\sup_{p \in B} \frac{\mathcal{L}(p_0)}{\mathcal{L}(p_*)} > e^{n\epsilon_n^2}\right) \\
&\equiv P_1 + P_2.
\end{aligned}$$

Now

$$P_1 \leq \sum_{p \in B} \mathbb{P}\left(\frac{\mathcal{L}(p)}{\mathcal{L}(p_*)} > e^{-n\epsilon_n^2 A^2/2}\right) \leq N(\epsilon/\sqrt{2C})e^{-n\epsilon_n^2 A^2/4} \leq e^{n\epsilon_n^2/(16C)}$$

where we used Lemma 36.72 and the definition of ϵ_n . To bound P_2 , define $K_n = \frac{1}{n} \sum_{i=1}^n \log \frac{p_0(Z_i)}{p_*(Z_i)}$.

Hence, $\mathbb{E}(K_n) = \text{KL}(p_0, p_*) \leq \epsilon^2/2$. Also,

$$\begin{aligned} \sigma^2 &\equiv \text{Var} \left(\log \frac{p_0(Z)}{p_*(Z)} \right) \leq \mathbb{E} \left(\log \frac{p_0(Z)}{p_*(Z)} \right)^2 \leq \log \left(\frac{c_2}{c_1} \right) \mathbb{E} \left(\log \frac{p_0(Z)}{p_*(Z)} \right) \\ &= \log \left(\frac{c_2}{c_1} \right) \text{KL}(p_0, p_*) \leq \log \left(\frac{c_2}{c_1} \right) \frac{\epsilon^2}{2} \equiv c_3 \epsilon_n^2 \end{aligned}$$

where we used (36.70). So, by Bernstein's inequality,

$$\begin{aligned} P_2 &= \mathbb{P}(K_n > \epsilon_n^2) = \mathbb{P}(K_n - \text{KL}(p_0, p_*) > \epsilon_n^2 - \text{KL}(p_0, p_*)) \\ &\leq \mathbb{P} \left(K_n - \text{KL}(p_0, p_*) > \frac{\epsilon_n^2}{2} \right) \leq 2 \exp \left(-\frac{n\epsilon_n^4}{8\sigma^2 + c_4\epsilon_n^2} \right) \\ &\leq 2 \exp(-c_5 n \epsilon_n^2). \end{aligned}$$

Thus, $P_1 + P_2 \leq \exp(-c_6 n \epsilon_n^2)$. Now,

$$\begin{aligned} \mathbb{E}(h(\hat{p}, p_0)) &= \int_0^{\sqrt{2}} \mathbb{P}(h(\hat{p}, p_0) > t) dt \\ &= \int_0^{D\epsilon_n} \mathbb{P}(h(\hat{p}, p_0) > t) dt + \int_{D\epsilon_n}^{\sqrt{2}} \mathbb{P}(h(\hat{p}, p_0) > t) dt \\ &\leq D\epsilon_n + \exp(-c_6 n \epsilon_n^2) \leq c_7 \epsilon_n. \end{aligned}$$

□

Lower Bound. Now we derive the lower bound.

36.74 Theorem. Let ϵ_n be the smallest ϵ such that $H(a\epsilon) \geq 64C^2 n \epsilon^2$. Then

$$\inf_{\hat{p}} \sup_{P \in \mathcal{P}} \mathbb{E}_p(h(p, \hat{p})) = \mathcal{X}(\epsilon_n).$$

Proof. Pick any $p \in \mathcal{P}$. Let $B = \{q : h(p, q) \leq 4\epsilon_n\}$. Let $F = \{p_1, \dots, p_N\}$ be an ϵ_n packing set for B . Then

$$N = \log P(\epsilon_n, B, h) \geq \log H(\epsilon_n, B, h) \geq \log H(a\epsilon_n) \geq 64C^2 n \epsilon_n^2.$$

Hence, for any $P_j, P_k \in F$,

$$\text{KL}(P_j^n, P_k^n) = n \text{KL}(P_j, P_k) \leq C n h^2(P_j, P_k) \leq 16C n \epsilon_n^2 \leq \frac{N}{4}.$$

It follows from Fano's inequality that

$$\inf_{\hat{p}} \sup_{p \in \mathcal{P}} \mathbb{E}_p h(p, \hat{p}) \geq \frac{1}{4} \min_{j \neq k} h(p_j, p_k) \geq \frac{\epsilon_n}{4}$$

as claimed. \square

In summary, we get the minimax rate by solving

$$H(\epsilon_n) \asymp n\epsilon_n^2.$$

Now we can use the Le Cam equation to compute some rates:

36.75 Example. Here are some standard examples:

Space	Entropy	Rate
Sobolev α	$\epsilon^{-1/\alpha}$	$n^{-\alpha/(2\alpha+1)}$
Sobolev α dimension d	$\epsilon^{-d/\alpha}$	$n^{-\alpha/(2\alpha+d)}$
Lipschitz α	$\epsilon^{-d/\alpha}$	$n^{-\alpha/(2\alpha+d)}$
Monotone	$1/\epsilon$	$n^{-1/3}$
Besov $B_{p,q}^\alpha$	$\epsilon^{-d/\alpha}$	$n^{-\alpha/(2\alpha+d)}$
Neural Nets	see below	see below
m -dimensional parametric model	$m \log(1/\epsilon)$	(m/n)

In the neural net case we have $f(x) = c_0 + \sum_i c_i \sigma(v_i^T x + b_i)$ where $\|c\|_1 \leq C$, $\|v_i\| = 1$ and σ is a step function or a Lipschitz sigmoidal function. Then

$$\left(\frac{1}{\epsilon}\right)^{1/2+1/d} \leq H(\epsilon) \leq \left(\frac{1}{\epsilon}\right)^{1/2+1/(2d)} \log(1/\epsilon) \tag{36.76}$$

and hence

$$n^{-(1+2/d)/(2+1/d)} (\log n)^{-(1+2/d)(1+1/d)/(2+1/d)} \leq \epsilon_n \leq (n/\log n)^{-(1+1/d)/(2+1/d)}. \tag{36.77}$$

\square

36.10 Summary

Minimax theory allows us to state precisely the best possible performance of any procedure under given conditions. The key tool for finding lower bounds on the minimax risk is Fano's inequality. Finding an upper bound usually involves finding a specific estimator and computing its risk.

36.11 Bibliographic remarks

There is a vast literature on minimax theory however much of it is scattered in various journal articles. Some texts that contain minimax theory include Tsybakov (2009), van de Geer (2000), van der Vaart (1998) and Wasserman (2006).

36.12 Appendix

36.12.1 Metrics For Probability Distributions

Minimax theory often makes use of various metrics for probability distributions. Here we summarize some of these metrics and their properties.

Let P and Q be two distributions with densities p and q . We write the distance between P and Q as either $d(P, Q)$ or $d(p, q)$ whichever is convenient. We define the following distances and related quantities.

Total variation	$\text{TV}(P, Q) = \sup_A P(A) - Q(A) $
L_1 distance	$d_1(P, Q) = \int p - q $
L_2 distance	$d_2(P, Q) = \sqrt{\int p - q ^2}$
Hellinger distance	$h(P, Q) = \sqrt{\int (\sqrt{p} - \sqrt{q})^2}$
Kullback-Leibler distance	$\text{KL}(P, Q) = \int p \log(p/q)$
χ^2	$\chi^2(P, Q) = \int (p - q)^2/p$
Affinity	$\ P \wedge Q\ = \int p \wedge q = \int \min\{p(x), q(x)\} dx$
Hellinger affinity	$A(P, Q) = \int \sqrt{pq}$

There are many relationships between these quantities. These are summarized in the next two theorems. We leave the proofs as exercises.

36.78 Theorem. *The following relationships hold:*

1. $\text{TV}(P, Q) = \frac{1}{2}d_1(P, Q) = 1 - \|p \wedge q\|$. (*Scheffé's Theorem.*)
2. $\text{TV}(P, Q) = P(A) - Q(A)$ where $A = \{x : p(x) > q(x)\}$.
3. $0 \leq h(P, Q) \leq \sqrt{2}$.
4. $h^2(P, Q) = 2(1 - A(P, Q))$.
5. $\|P \wedge Q\| = 1 - \frac{1}{2}d_1(P, Q)$.
6. $\|P \wedge Q\| \geq \frac{1}{2}A^2(P, Q) = \frac{1}{2} \left(1 - \frac{h^2(P, Q)}{2}\right)^2$. (*Le Cam's inequalities.*)
7. $\frac{1}{2}h^2(P, Q) \leq \text{TV}(P, Q) = \frac{1}{2}d_1(P, Q) \leq h(P, Q) \sqrt{1 - \frac{h^2(P, Q)}{4}}$.
8. $\text{TV}(P, Q) \leq \sqrt{\text{KL}(P, Q)}/2$. (*Pinsker's inequality.*)
9. $\|P \wedge Q\| \geq \frac{1}{2}e^{-\text{KL}(P, Q)}$.
10. $\text{TV}(P, Q) \leq h(P, Q) \leq \sqrt{\text{KL}(P, Q)} \leq \sqrt{\chi^2(P, Q)}$.

Let P^n denote the product measure based on n independent samples from P .

36.79 Theorem. *The following relationships hold:*

1. $h^2(P^n, Q^n) = 2 \left(1 - \left(1 - \frac{h^2(P, Q)}{2}\right)^n\right)$.

2. $\|P^n \wedge Q^n\| \geq \frac{1}{2} A^2(P^n, Q^n) = \frac{1}{2} (1 - \frac{1}{2} h^2(P, Q))^{2n}$.
3. $\|P^n \wedge Q^n\| \geq (1 - \frac{1}{2} d_1(P, Q))^n$.
4. $\text{KL}(P^n, Q^n) = n \text{KL}(P, Q)$.

36.12.2 Fano's Lemma

For $0 < p < 1$ define the entropy $h(p) = -p \log p - (1-p) \log(1-p)$ and note that $0 \leq h(p) \leq \log 2$. Let (Y, Z) be a pair of random variables each taking values in $\{1, \dots, N\}$ with joint distribution $P_{Y,Z}$. Then the mutual information is defined to be

$$I(Y; Z) = \text{KL}(P_{Y,Z}, P_Y \times P_Z) = H(Y) - H(Y | Z) \quad (36.80)$$

where $H(Y) = -\sum_j \mathbb{P}(Y = j) \log \mathbb{P}(Y = j)$ is the entropy of Y and $H(Y | Z)$ is the entropy of Y given Z . We will use the fact that $I(Y; h(Z)) \leq I(Y; Z)$ for an function h .

36.81 Lemma. *Let Y be a random variable taking values in $\{1, \dots, N\}$. Let $\{P_1, \dots, P_N\}$ be a set of distributions. Let X be drawn from P_j for some $j \in \{1, \dots, N\}$. Thus $P(X \in A | Y = j) = P_j(A)$. Let $Z = g(X)$ be an estimate of Y taking values in $\{1, \dots, N\}$. Then,*

$$H(Y | X) \leq \mathbb{P}(Z \neq Y) \log(N-1) + h(\mathbb{P}(Z = Y)). \quad (36.82)$$

We follow the proof from Cover and Thomas (1991).

Proof. Let $E = I(Z \neq Y)$. Then

$$H(E, Y | X) = H(Y | X) + H(E | X, Y) = H(Y | X)$$

since $H(E | X, Y) = 0$. Also,

$$H(E, Y | X) = H(E | X) + H(Y | E, X).$$

But $H(E | X) \leq H(E) = h(P(Z = Y))$. Also,

$$\begin{aligned} H(Y | E, X) &= P(E = 0)H(Y | X, E = 0) + P(E = 1)H(Y | X, E = 1) \\ &\leq P(E = 0) \times 0 + h(P(Z = Y)) \log(N-1) \end{aligned}$$

since $Y = g(X)$ when $E = 0$ and, when $E = 1$, $H(Y | X, E = 1)$ by the log number of remaining outcomes. Combining these gives $H(Y | X) \leq \mathbb{P}(Z \neq Y) \log(N-1) + h(\mathbb{P}(Z = Y))$. \square

36.83 Lemma. (Fano's Inequality) Let $\mathcal{P} = \{P_1, \dots, P_N\}$ and $\beta = \max_{j \neq k} \text{KL}(P_j, P_k)$. For any random variable Z taking values on $\{1, \dots, N\}$,

$$\frac{1}{N} \sum_{j=1}^N P_j(Z \neq j) \geq \left(1 - \frac{n\beta + \log 2}{\log N}\right).$$

Proof. For simplicity, assume that $n = 1$. The general case follows since $\text{KL}(P^n, Q^n) = n\text{KL}(P, Q)$. Let Y have a uniform distribution on $\{1, \dots, N\}$. Given $Y = j$, let X have distribution P_j . This defines a joint distribution P for (X, Y) given by

$$P(X \in A, Y = j) = P(X \in A | Y = j)P(Y = j) = \frac{1}{N}P_j(A).$$

Hence,

$$\frac{1}{N} \sum_{j=1}^N P(Z \neq j | Y = j) = P(Z \neq Y).$$

From (36.82),

$$\begin{aligned} H(Y | Z) &\leq P(Z \neq Y) \log(N-1) + h(P(Z = Y)) \leq P(Z \neq Y) \log(N-1) + h(1/2) \\ &= P(Z \neq Y) \log(N-1) + \log 2. \end{aligned}$$

Therefore,

$$\begin{aligned} P(Z \neq Y) \log(N-1) &\geq H(Y | Z) - \log 2 = H(Y) - I(Y; Z) - \log 2 \\ &= \log N - I(Y; Z) - \log 2 \geq \log N - \beta - \log 2. \end{aligned} \quad (36.84)$$

The last inequality follows since

$$I(Y; Z) \leq I(Y; X) = \frac{1}{N} \sum_{j=1}^N \text{KL}(P_j, \bar{P}) \leq \frac{1}{N^2} \sum_{j,k} \text{KL}(P_j, P_k) \leq \beta \quad (36.85)$$

where $\bar{P} = N^{-1} \sum_{j=1}^N P_j$ and we used the convexity of K . Equation (36.84) shows that

$$P(Z \neq Y) \log(N-1) \geq \log N - \beta - \log 2$$

and the result follows. \square

36.12.3 Assouad's Lemma

Assouad's Lemma is another way to get a lower bound using hypercubes. Let

$$\mathcal{X} = \left\{ \omega = (\omega_1, \dots, \omega_N) : \omega_j \in \{0, 1\} \right\}$$

be the set of binary sequences of length N . Let $\mathcal{P} = \{P_\omega : \omega \in \mathcal{X}\}$ be a set of 2^N distributions indexed by the elements of \mathcal{X} . Let $h(\omega, \nu) = \sum_{j=1}^N I(\omega_j \neq \nu_j)$ be the *Hamming distance* between $\omega, \nu \in \mathcal{X}$.

36.86 Lemma. Let $\{P_\omega : \omega \in \mathcal{X}\}$ be a set of distributions indexed by ω and let $\theta(P)$ be a parameter. For any $p > 0$ and any metric d ,

$$\max_{\omega \in \mathcal{X}} E_\omega \left(d^p(\hat{\theta}, \theta(P_\omega)) \right) \geq \frac{N}{2^{p+1}} \left(\min_{\substack{\omega, \nu \\ h(\omega, \nu) \neq 0}} \frac{d^p(\theta(P_\omega), \theta(P_\nu))}{h(\omega, \nu)} \right) \left(\min_{\substack{\omega, \nu \\ h(\omega, \nu) = 1}} \|P_\omega \wedge P_\nu\| \right). \quad (36.87)$$

For a proof, see van der Vaart (1998) or Tsybakov (2009).

Exercises

- 36.1 Let $P = N(\mu_1, \Sigma_1)$ and $Q = N(\mu_2, \Sigma_2)$ be two, d -dimensional multivariate Gaussian distributions. Find $\text{KL}(P, Q)$.
- 36.2 Prove Lemma 36.16.
- 36.3 Prove Lemma 36.17.
- 36.4 Prove Theorem 36.78.
- 36.5 Prove Theorem 36.79.
- 36.6 Prove (36.80).
- 36.7 Prove that $I(Y; h(Z)) \leq I(Y; Z)$.
- 36.8 Prove (36.85).
- 36.9 Prove (36.39).
- 36.10 Prove Lemma ??.
- 36.11 Show that (36.40) holds.
- 36.12 Construct F so that (36.41) holds.
- 36.13 Verify the result in Example 36.63.
- 36.14 Verify the results in Example 36.66.
- 36.15 Prove Lemma 36.62.
- 36.16 Prove equation ??.
- 36.17 Prove equation ??.
- 36.18 Prove equation ??.