

Lecture 14: September 30

Lecturer: Siva Balakrishnan

Today we will discuss *point estimation*, i.e. the task is given $X_1, \dots, X_n \sim p(X; \theta)$ we would like to construct an *estimator* $\hat{\theta}(X_1, \dots, X_n)$. The rationale for the task is simple – it is closely related to the central preoccupation of Statistics/ML – we would like to understand the population given a random sample, i.e. we want to understand $p(X; \theta)$ and one way to do this is to estimate θ . In some cases, θ has some natural meaning (maybe it is the mean of the population) but this is not necessarily the case.

In the next few lectures we will discuss ways to construct estimators and then how to compare or evaluate them. Roughly, the questions we are trying to answer are:

1. Are there general purpose methods to come up with estimators of θ ?
2. Given two (or more) estimators is there a general framework in which we can compare estimators?
3. Finally, are there general purpose ways to analyze complex estimators (say estimators that are not simple averages)?

As a small point – we usually use “estimator” to refer to a random variable (a statistic, a function of the sample) and “estimate” to refer to its realized value. We have already studied estimation in a relatively simple context in great detail, given a sample, we aimed at estimating the mean of the population. When θ is not a simple mean then we need to think a bit harder to decide how to estimate it. We will focus on general purpose methods for estimation.

$X_1, \dots, X_n \sim p(x; \theta)$. Want to estimate $\theta = (\theta_1, \dots, \theta_k)$. An *estimator*

$$\hat{\theta} = \hat{\theta}_n = w(X_1, \dots, X_n)$$

is a function of the data. Keep in mind that the parameter is a fixed, unknown constant. The estimator is a random variable.

For now, we will discuss three methods of constructing estimators:

1. The Method of Moments (MOM)
2. Maximum likelihood (MLE)
3. Bayesian estimators.

14.1 The Method of Moments

Suppose that $\theta = (\theta_1, \dots, \theta_k)$. Define

$$\begin{aligned} m_1 &= \frac{1}{n} \sum_{i=1}^n X_i, & \mu_1(\theta) &= \mathbb{E}(X_i) \\ m_2 &= \frac{1}{n} \sum_{i=1}^n X_i^2, & \mu_2(\theta) &= \mathbb{E}(X_i^2) \\ & & \vdots & \\ m_k &= \frac{1}{n} \sum_{i=1}^n X_i^k, & \mu_k(\theta) &= \mathbb{E}(X_i^k). \end{aligned}$$

Let $\hat{\theta} = (\hat{\theta}_1, \dots, \hat{\theta}_k)$ solve:

$$m_j = \mu_j(\hat{\theta}), \quad j = 1, \dots, k.$$

In other words, we equate the first k sample moments with the first k theoretical moments. This defines k equations with k unknowns.

Example 14.1 $N(\beta, \sigma^2)$ with $\theta = (\beta, \sigma^2)$. Then $\mu_1 = \beta$ and $\mu_2 = \sigma^2 + \beta^2$. Equate:

$$\frac{1}{n} \sum_{i=1}^n X_i = \hat{\beta}, \quad \frac{1}{n} \sum_{i=1}^n X_i^2 = \hat{\sigma}^2 + \hat{\beta}^2$$

to get

$$\hat{\beta} = \bar{X}_n, \quad \hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \bar{X}_n)^2.$$

Example 14.2 Suppose

$$X_1, \dots, X_n \sim \text{Binomial}(k, p)$$

where both k and p are unknown. We get

$$kp = \bar{X}_n, \quad \frac{1}{n} \sum_{i=1}^n X_i^2 = kp(1-p) + k^2p^2$$

giving

$$\hat{p} = \frac{\bar{X}_n}{k}, \quad \hat{k} = \frac{\bar{X}_n^2}{\bar{X}_n - \frac{1}{n} \sum_{i=1}^n (X_i - \bar{X}_n)^2}.$$

The method of moments was popular many years ago because it is often easy to compute. Lately, it has attracted attention again. For example, there is a large literature on estimating “mixtures of Gaussians” using the method of moments.

14.2 Maximum Likelihood

The most popular method for estimating parameters is maximum likelihood. One of the reasons is that, under certain conditions, the maximum likelihood estimator is optimal. We'll discuss optimality later.

The maximum likelihood estimator (mle) $\hat{\theta}$ is defined as the maximizer of

$$\mathcal{L}(\theta) = p(X_1, \dots, X_n; \theta) \stackrel{iid}{=} \prod_i p(X_i; \theta).$$

This is the same as maximizing the log-likelihood

$$\mathcal{L}\mathcal{L}(\theta) = \log \mathcal{L}(\theta).$$

Often it suffices to solve

$$\frac{\partial \mathcal{L}\mathcal{L}(\theta)}{\partial \theta_j} = 0, \quad j = 1, \dots, k.$$

Example 14.3 *Binomial.* $\mathcal{L}(p) = \prod_i p^{X_i} (1-p)^{1-X_i} = p^S (1-p)^{n-S}$ where $S = \sum_i X_i$. So

$$\mathcal{L}\mathcal{L}(p) = S \log p + (n-S) \log(1-p)$$

and $\hat{p} = \bar{X}_n$.

Example 14.4 $X_1, \dots, X_n \sim N(\mu, 1)$.

$$\mathcal{L}(\mu) \propto \prod_i e^{-(X_i - \mu)^2/2} \propto e^{-n(\bar{X}_n - \mu)^2}, \quad \mathcal{L}\mathcal{L}(\mu) = -\frac{n}{2}(\bar{X}_n - \mu)^2$$

and $\hat{\mu} = \bar{X}_n$. For $N(\mu, \sigma^2)$ we have

$$\mathcal{L}(\mu, \sigma^2) \propto \prod_i \frac{1}{\sigma} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{i=1}^n (X_i - \mu)^2 \right\}$$

and

$$\mathcal{L}\mathcal{L}(\mu, \sigma^2) = -n \log \sigma - \frac{1}{2\sigma^2} \sum_{i=1}^n (X_i - \mu)^2.$$

Set

$$\frac{\partial \mathcal{L}}{\partial \mu} = 0, \quad \frac{\partial \mathcal{L}}{\partial \sigma^2} = 0$$

to get

$$\hat{\mu} = \frac{1}{n} \sum_{i=1}^n X_i, \quad \hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \bar{X})^2.$$

Example 14.5 Let $X_1, \dots, X_n \sim \text{Uniform}(0, \theta)$. Then

$$\mathcal{L}(\theta) = \frac{1}{\theta^n} I(\theta > X_{(n)})$$

and so $\hat{\theta} = X_{(n)}$.

What is the method of moments estimator above? Can you do a back of the envelope computation to compare the two estimators?

14.2.1 MLE and MoM for exponential families

The log-likelihood in an exponential family is concave and given by

$$\mathcal{L}(\theta; x_1, \dots, x_n) \propto \left[\sum_{i=1}^s \theta_i \sum_{j=1}^n T_i(x_j) - nA(\theta) \right],$$

so we can simply take the derivative with respect to θ and set this equal to 0. Using the facts we have seen in the last lecture about the derivative of A , we can see that this amounts to solving the following system of equations for θ :

$$\mathbb{E}_{p(X;\theta)}[T_i(X)] = \frac{1}{n} \sum_{j=1}^n T_i(x_j) \quad \text{for } i \in \{1, \dots, s\}.$$

So the maximum likelihood estimator simply picks the parameters θ to match the empirical expectations of the sufficient statistics to the expected value of the sufficient statistics under the distribution.

Usually we cannot compute this estimator in closed form so we use an iterative algorithm (like gradient ascent) to maximize the likelihood. However, you should remember that exponential families have concave likelihoods so this is usually a tractable endeavour (at least for simple enough families).

For exponential families as we can see above the method of moments coincides with the MLE (if we chose the sufficient statistics to direct which moments to compute).

14.2.2 Equivariance and the profile likelihood

Suppose that $\theta = (\eta, \xi)$. The *profile likelihood* for η is defined by

$$\mathcal{L}(\eta) = \sup_{\xi} \mathcal{L}(\eta, \xi).$$

To find the mle of η we can proceed in two ways. We could find the overall mle $\hat{\theta} = (\hat{\eta}, \hat{\xi})$. The mle for η is just the first coordinate of $(\hat{\eta}, \hat{\xi})$. Alternatively, we could find the maximizer of the profile likelihood. These give the same answer. Do you see why?

The mle is *equivariant*. if $\eta = g(\theta)$ then $\hat{\eta} = g(\hat{\theta})$. Suppose g is invertible so $\eta = g(\theta)$ and $\theta = g^{-1}(\eta)$. Define $\mathcal{L}^*(\eta) = \mathcal{L}(\theta)$ where $\theta = g^{-1}(\eta)$. So, for any η ,

$$\mathcal{L}^*(\hat{\eta}) = \mathcal{L}(\hat{\theta}) \geq \mathcal{L}(\theta) = \mathcal{L}^*(\eta)$$

and hence $\hat{\eta} = g(\hat{\theta})$ maximizes $\mathcal{L}^*(\eta)$. For non invertible functions this is still true if we define

$$\mathcal{L}^*(\eta) = \sup_{\theta: g(\theta)=\eta} \mathcal{L}(\theta).$$

(In other words, the profile likelihood.)

Example 14.6 *Binomial.* The mle is $\hat{p} = \bar{X}_n$. Let $\psi = \log(p/(1-p))$. Then $\hat{\psi} = \log(\hat{p}/(1-\hat{p}))$.

14.3 Bayes Estimator

To define the Bayes estimator, we begin by treating θ as a random variable. This point requires much discussion (which we will have later). For now, just tentatively think of θ as random. We start with a *prior distribution* $p(\theta)$ on θ . Note that

$$p(x_1, \dots, x_n | \theta) p(\theta) = p(x_1, \dots, x_n, \theta).$$

Now compute the *posterior distribution* by Bayes' theorem:

$$p(\theta | x_1, \dots, x_n) = \frac{p(x_1, \dots, x_n | \theta) p(\theta)}{p(x_1, \dots, x_n)}$$

where

$$p(x_1, \dots, x_n) = \int p(x_1, \dots, x_n | \theta) p(\theta) d\theta.$$

This can be written as

$$p(\theta|x_1, \dots, x_n) \propto \mathcal{L}(\theta)p(\theta) = \text{Likelihood} \times \text{prior}.$$

Now compute a point estimator from the posterior. For example:

$$\hat{\theta} = \mathbb{E}(\theta|x_1, \dots, x_n) = \int \theta p(\theta|x_1, \dots, x_n) d\theta = \frac{\int \theta p(x_1, \dots, x_n|\theta) p(\theta) d\theta}{\int p(x_1, \dots, x_n|\theta) p(\theta) d\theta}.$$

Example 14.7 Let $X_1, \dots, X_n \sim \text{Bernoulli}(\theta)$. Let the prior be $\theta \sim \text{Beta}(\alpha, \beta)$. Hence

$$p(\theta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1} (1 - \theta)^{\beta-1},$$

and

$$\Gamma(\alpha) = \int_0^\infty t^{\alpha-1} e^{-t} dt.$$

Set $Y = \sum_i X_i$. Then

$$p(\theta|X) \propto \underbrace{\theta^Y (1 - \theta)^{n-Y}}_{\text{likelihood}} \times \underbrace{\theta^{\alpha-1} (1 - \theta)^{\beta-1}}_{\text{prior}} \propto \theta^{Y+\alpha-1} (1 - \theta)^{n-Y+\beta-1}.$$

Therefore, $\theta|X \sim \text{Beta}(Y + \alpha, n - Y + \beta)$. The Bayes estimator is

$$\tilde{\theta} = \frac{Y + \alpha}{(Y + \alpha) + (n - Y + \beta)} = \frac{Y + \alpha}{\alpha + \beta + n} = (1 - \lambda)\hat{\theta}_{mle} + \lambda \bar{\theta}$$

where

$$\bar{\theta} = \frac{\alpha}{\alpha + \beta}, \quad \lambda = \frac{\alpha + \beta}{\alpha + \beta + n}.$$

This is an example of a conjugate prior.

Example 14.8 Let $X_1, \dots, X_n \sim N(\mu, \sigma^2)$ with σ^2 known. Let $\mu \sim N(m, \tau^2)$. Then

$$\mathbb{E}(\mu|X) = \frac{\tau^2}{\tau^2 + \frac{\sigma^2}{n}} \bar{X}_n + \frac{\frac{\sigma^2}{n}}{\tau^2 + \frac{\sigma^2}{n}} m$$

and

$$\text{Var}(\mu|X) = \frac{\sigma^2 \tau^2 / n}{\tau^2 + \frac{\sigma^2}{n}}.$$