

MS&E 226: “Small” Data

Lecture 11: Maximum likelihood (v2)

Ramesh Johari
ramesh.johari@stanford.edu

The likelihood function

Estimating the parameter

This lecture develops the methodology behind the *maximum likelihood* approach to parameter estimation.

Basic idea:

- ▶ In guessing what the parameters are, we first ask: what is the chance of seeing the data we observe, given a particular value of the parameter?
- ▶ We then pick the parameter values that maximize this chance.

Example 1: Flipping a biased coin

Suppose we flip a coin n times, and get observations $\mathbf{Y} = (Y_1, \dots, Y_n)$.

If the bias on the coin was q , what is the probability we get data \mathbf{Y} ?

This is the *likelihood* of \mathbf{Y} given q :

$$f(\mathbf{Y}|q) = \prod_{i=1}^n q^{Y_i} (1 - q)^{1 - Y_i}.$$

Example 2: Linear normal population model

Suppose that you are given p -dimensional covariate vectors \mathbf{X}_i , $i = 1, \dots, n$, together with corresponding observations Y_i , $i = 1, \dots, n$.

Suppose you assume the *linear normal population model* with *i.i.d. errors*:

$$Y_i = \mathbf{X}_i\boldsymbol{\beta} + \varepsilon_i,$$

where $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$, and the ε_i are i.i.d.

What are the parameters? What is the likelihood?

Example 2: Linear normal population model

Parameters are β and σ^2 .

In deriving likelihood, we will typically treat \mathbf{X} as *given*; i.e., we focus on inference about the relationship between \mathbf{X} and \mathbf{Y} , rather than the distribution of \mathbf{X} .

Likelihood is probability density of seeing \mathbf{Y} , given parameters and \mathbf{X} :

$$f(\mathbf{Y}|\beta, \sigma^2, \mathbf{X}) = \prod_{i=1}^n \left(\frac{1}{\sqrt{2\pi\sigma^2}} \right) \exp \left(-\frac{(Y_i - \mathbf{X}_i\beta)^2}{2\sigma^2} \right).$$

Example 2: Linear normal population model

Derivation of the likelihood:

In general: Parametric likelihood

In general, suppose there is some process that generates data \mathbf{Y} .¹

Suppose each choice of a parameter vector θ gives rise to a conditional pmf (or pdf) $f(\mathbf{Y}|\theta)$.

This is the probability (or density) of seeing \mathbf{Y} , given parameters θ .

We call this the *likelihood* of \mathbf{Y} given θ .

¹As noted on the previous slide, in the case of regression, we treat \mathbf{X} as given and look at the process that generates \mathbf{Y} from \mathbf{X} .

Log likelihood

It is often easier to work with logs when taking likelihoods (not least because multiplying many small probabilities together can cause numerical instability).

We call this the *log likelihood function* (LLF).

Example 1 (biased coin):

$$\log f(\mathbf{Y}|q) = \sum_{i=1}^n Y_i \log q + (1 - Y_i) \log(1 - q).$$

Example 2 (linear normal population model):

$$\log f(\mathbf{Y}|\boldsymbol{\beta}, \sigma^2, \mathbf{X}) = -\frac{n}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} \sum_{i=1}^n (Y_i - \mathbf{X}_i\boldsymbol{\beta})^2.$$

Maximum likelihood

Maximizing the log likelihood

The *maximum likelihood estimate* (MLE) is the parameter value that maximizes the likelihood.

It is found by solving the following optimization problem:

$$\begin{array}{ll} \text{maximize} & f(\mathbf{Y}|\boldsymbol{\theta}) \quad (\text{or } \log f(\mathbf{Y}|\boldsymbol{\theta})) \\ & \text{over feasible choices of } \boldsymbol{\theta}. \end{array}$$

We denote the resulting solution by $\hat{\boldsymbol{\theta}}_{\text{MLE}}$.

Example 1: Flipping a biased coin

Suppose we maximize the log likelihood. The resulting solution is:

$$\hat{q}_{\text{MLE}} = \frac{1}{n} \sum_{i=1}^n Y_i = \bar{Y}.$$

This is fairly reasonable: we estimate q by the mean number of successes.

Example 1: Flipping a biased coin

Derivation of the MLE estimate:

Example 2: Linear normal model, known σ^2

Suppose that σ^2 is known, so the goal is to estimate β . Returning to LLF, our problem is equivalent to choosing $\hat{\beta}$ to minimize:

$$\text{minimize } \sum_{i=1}^n (Y_i - \mathbf{X}_i \hat{\beta})^2.$$

In other words: the OLS solution is the MLE estimate of the coefficients!

Example 2: Linear normal model, known σ^2

Let's pause to reflect on what we've found here.

When we first discussed OLS, we did so from a purely algebraic view — no assumptions.

Example 2: Linear normal model, known σ^2

Let's pause to reflect on what we've found here.

When we first discussed OLS, we did so from a purely algebraic view — no assumptions.

When we discussed prediction, we showed that as long as the population model is linear and the ε_i are *uncorrelated* with \mathbf{X}_i , then OLS is unbiased, and has minimum variance among unbiased linear estimators (Gauss-Markov theorem).

Example 2: Linear normal model, known σ^2

Let's pause to reflect on what we've found here.

When we first discussed OLS, we did so from a purely algebraic view — no assumptions.

When we discussed prediction, we showed that as long as the population model is linear and the ε_i are *uncorrelated* with \mathbf{X}_i , then OLS is unbiased, and has minimum variance among unbiased linear estimators (Gauss-Markov theorem).

Now, we have shown that in addition if we assume the ε_i are i.i.d. normal random variables, then OLS is the maximum likelihood estimate.

Example 2: Linear normal model, unknown σ^2

What happens if σ^2 is *unknown*?

The MLE for β remains unchanged (the OLS solution), and the MLE estimate for σ^2 is:

$$\hat{\sigma}_{\text{MLE}}^2 = \frac{1}{n} \sum_{i=1}^n (Y_i - \mathbf{X}_i \hat{\beta})^2 = \frac{1}{n} \sum_{i=1}^n r_i^2.$$

This is intuitive: the sum of squared residuals is an estimate of the variance of the error.

Example 3: Logistic regression

The logistic regression parameters are found through *maximum likelihood*.

The (conditional) likelihood for \mathbf{Y} given \mathbf{X} and parameters β is:

$$\mathbb{P}(\mathbf{Y}|\beta, \mathbf{X}) = \prod_{i=1}^n g^{-1}(\mathbf{X}_i\beta)^{Y_i} (1 - g^{-1}(\mathbf{X}_i\beta))^{1-Y_i}$$

where $g^{-1}(z) = \exp(z)/(1 + \exp(z))$.

Let $\hat{\beta}_{\text{MLE}}$ be the resulting solution; these are the *logistic regression coefficients*.

Example 3: Logistic regression [*]

Unfortunately, in contrast to our previous examples, maximum likelihood estimation does not have a closed form solution in the case of logistic regression.

However, it turns out that there are reasonably efficient iterative methods for algorithmically computing the MLE solution.

One example is an algorithm inspired by weighted least squares, called *iteratively reweighted least squares*. This algorithm iteratively updates the weights in WLS to converge to the logistic regression MLE solution. See [AoS], Section 13.7 for details.