Maximum Likelihood Estimation

November 16, 2023

Maximum Likelihood Estimation



It is a part of probability that many improbable things will happen.

Aristotle

Maximum Likelihood Estimation

Estimation Problem

- Consider a sample X that comes from an unknown member P_{θ0} of a family of distributions {P_θ}_{θ∈Θ}.
- Any statistic $\hat{\theta}$ that is computable from the observed values **x** of **X** and aims at approximating true θ_0 can be called an estimator of this parameter.
- We are interested only in 'good estimators'.
- For example, we want the estimator error $\epsilon = \hat{\theta} \theta_0$ to be small, for example its MSE to be small

$$\mathsf{E}_{ heta_0}(\epsilon^2) = \mathsf{Var}_{ heta_0}(\hat{ heta}) + \left(\mathsf{E}_{ heta_0}\hat{ heta} - heta_0
ight)^2 = \mathsf{Var}_{ heta_0}(\hat{ heta}) + \mathsf{Bias}_{ heta_0}^2.$$

 Ideally we would like to have the bias to be zero and the variance to be minimal.

Estimator maximizing the likelihood

- *L*_x(θ) the likelihood function of θ (the probability density depending on θ with given x).
- We think of *L*_x(θ) as a measure of how "likely" θ is to have produced the observed x.
- The method of maximum likelihood consists of finding that value $\hat{\theta}_{MLE}$ of the parameter that is "most likely" to have produced the data.
- That is, if **X** = **x** is observed

$$\hat{\theta}_{MLE} = \operatorname*{argmax}_{\theta \in \Theta} L_{\mathbf{x}}(\theta).$$

- Maximum likelihood estimates need neither exist nor be unique.
- But in a wide important case, they are asymptotically (with respect to the sample size) the best possible ones.

Two examples

An iid sample from normal distribution with θ = (μ, σ²). It is easy to check that

$$\hat{ heta} = (ar{X}, ar{X^2} - ar{X}^2)$$

is the MLE for $\boldsymbol{\theta}$

An iid sample from uniform distribution on a sequence {1,...,θ} with the size of the population θ ∈ N being an unknown parameter. It is easy to see that the likelihood has the form

$$L_{\mathbf{x}}(\theta) = \frac{1}{\theta^n} \mathbf{1}_{[0,\theta]}(\max_{i=1,\dots,n} x_i)$$

and the maximum is reached at $\hat{\theta} = \max_{i=1,...,n} x_i$. Similar argument leads to the same estimator if the sample is from the continuous uniform distribution on $[0, \theta]$.

Likelihood equations

• Suppose:

- $X \sim P_{\theta}$, with $\theta \in \Theta$, an open set parameter space
- the likelihood function $L_X(\theta)$ is differentiable in θ
- $\hat{\theta}_{MLE}$ exists

Then: $\hat{\theta}_{MLE}$ must satisfy the Likelihood Equation(s)

$$abla_{ heta} L_X(heta) = \mathbf{0}.$$

Important Case: for independent X_i;s with

$$\sum_{i=1}^n \nabla \log p_i(x_i|\theta) = 0$$

NOTE: $p_i(\cdot|\theta)$ may vary with *i*.

Multinomial trials

- an experiment with n i.i.d. trials in which each trial can produce a result in one of k categories
- $X_i = j$ if the *i*th trial produces a result in the *j*th category
- let $\theta_j = P(X_i = j)$ be the probability of the *j*th category, and let

$$N_j = \sum_{i=1}^n \mathbf{1}_{\{X_i=j\}}$$

be the number of observations in the *j*th category.

• for an experiment in which we observe n_j , j = 1, ..., k we have

$$l_{\mathbf{x}}(\theta) = \sum_{j=1}^{k} n_j \log \theta_j,$$

with $\theta_k = 1 - \sum_{j=1}^{k-1} \theta_j$ and the normal (likelihood) equations are for j < k

$$\frac{\partial l_{\mathbf{x}}}{\partial \theta_j} = n_j / \theta_j - n_k / \left(1 - \sum_{i=1}^{k-1} \theta_i \right) = 0.$$

Thus, assuming w.l.g. that $n_k \neq 0$, $n_j/n_k = \theta_j/\theta_k$, which implies that $\hat{\theta}_j = n_j/n$.

Natural parameter

Questions of existence and uniqueness of maximum likelihood estimates in canonical exponential families can be answered completely and elegantly. This is largely a consequence of the strict concavity of the log likelihood in the natural parameter η .

Definition $\{P_{\theta}\}, \theta \in \Theta \text{ is a } k\text{-parameter exponential family if}$ $p(x|\theta) = h(x) \exp\left(\sum_{j=1}^{k} \eta_{j}(\theta) T_{j}(x) - B(\theta)\right), x \in \mathbb{R}^{q}$ • $\eta_{1}, \dots, \eta_{k}$ and B are real-valued functions mapping $\Theta \mapsto \mathbb{R}$, • T_{1}, \dots, T_{k} and h are real-valued functions mapping $\mathbb{R}^{q} \mapsto \mathbb{R}$.

Canonical form of exponential families

Canonical parameter

$$q(x,\eta) = h(x) \exp(\mathbf{T}^{\top}(x)\eta - A(\eta)).$$

In the continuous case

$$A(\eta) = \log \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} h(x) \exp(\mathbf{T}^{\top}(x)\eta) dx.$$

In the discrete case, $A(\eta)$ is defined in the same way except integrals are replaced by sums. In either case, we define the natural parameter space as

$$\mathcal{E} = \{\eta \in \mathbf{R}^k : -\infty < \mathcal{A}(\eta) < \infty\}.$$

Evidently, every *k*-dimensional parameter of an exponential family is also k'-dimensional with k' > k. However, there is a minimal dimension. An exponential family is of rank *k* iff the generating statistic **T** is *k*-dimensional and 1, $T_1(X), \ldots, T_k(X)$ are linearly independent with positive probability.

MLE in the convexity (concavity) context

Proposition 2.3.1.

Suppose that the domain of the parameters is open in \mathbf{R}^{p} . If the loglikelihood is the strictly concave function of the parameter and is unbounded on the boundary of the parameter space, then the MLE exists and is unique.

Some clarificiations:

- an open set is a set that with each point contains also a disk (ball) that is in this set – topological definition
- the boundary of a set U is a set ∂U that in any disk (ball) contains both the points from the set U and outside of the set U
- the strictly concave is a function like $-x^2$, ('sad face')

A necessary and sufficient condition for existence and uniqueness of the MLE

Theorem 2.3.1.

Suppose \mathcal{P} is the canonical exponential family generated by (\mathbf{T}, h) and that

- the natural parameter space, \mathcal{E} , is open,
- the family is of rank k.

Then the MLE of η exists and is unique for every *x* such that $t_0 = T(x)$ satisfies for all $\mathbf{c} \neq \mathbf{0}$:

$$\mathsf{P}(\mathbf{c}^{\top}\mathbf{T}(X) > \mathbf{c}^{\top}\mathbf{t}_{0}) > 0.$$
(1)

and the MLE is a solution to

$$\dot{A}(\eta)(=E_{\eta}(\mathbf{T}(X)))=\mathbf{t}_{0}.$$

Conversely, if t_0 doesn't satisfy (1), then the MLE doesn't exist and the above equation has no solution.

Example – the two-parameter gamma family

The density proportional to $\exp(-\lambda x)x^{p-1}$. This is a rank 2 canonical exponential family with

$$\mathbf{T} = \left(\sum \log X_i, \sum X_i\right)$$

$$h(x) = 1/x$$

$$\eta = (p, -\lambda)$$

$$A(\eta) = n \left(\log \Gamma(\eta_1) - \eta_1 \log(-\eta_2)\right).$$

We have

$$\dot{A}(\eta) = n(\Gamma'(p)/\Gamma(p) - \log \lambda, p/\lambda) = n(\overline{\log X}, \overline{X}).$$

We conclude from Theorem 2.3.2 that the equations have a unique solution with probability 1. How to find such nonexplicit solutions is discussed in Section 2.4 to which we turn next.

MLE as an optimization problem

- As we have seen, even in the context of canonical multiparameter exponential families, such as the two-parameter gamma, MLEs may not be given explicitly by formulae but only implicitly as the solutions of systems of nonlinear equations.
- Even in the classical regression model with design matrix of full rank, the formula for the parameter estimator is easy to write down symbolically but not easy to evaluate if the dimension of the parameter is large.
- It is because inversion of on the order of the third power operations (with respect to the parameter dimension) to evaluate the inversion matrix.
- We will discuss three algorithms of a type used in different statistical contexts both for their own sakes and to illustrate what kinds of things can be established about the black boxes to which we all, at various times, entrust ourselves.

Bisection algorithm

By the calculus, if the likelihood is differentiable and the location of the global maximum is in the interior of the parameter space finding the solution reduces to finding the solution to the likelihood (normal) equations (the equations for zeros of the derivative functions).

- Finding two points, one that the derivative is below zero and one that it is above.
- Take the middle of the interval made of these points and cut it in half.
- Evaluate the value of the derivative at the middle, and replace by the middle point the one of the original points that has the same sign of the value of the derivative.
- Continue until the length of the interval has the desired accuracy or the zero is reached before that.

Bisection for the one-dimensional exponential family

- The bisection algorithm applied to the strictly increasing derivative (concave function) finds the solution that is unique.
- For the canonical exponential family this is the case since $f(\eta) = E_{\eta}T(X) t_0$ is strictly increasing:

$$f'(\eta) = A''(\eta) = Var_{\eta}T(X) > 0.$$

• **Example:** For the shape parameter gamma family with the density proportional to $x^{\theta-1}e^{-x}$, one has to find the solution

$$\frac{\Gamma'(\theta)}{\Gamma(\theta)} = \overline{\log X}.$$

Extension to multidimensional parameter – coordinate ascent

We need to find a solution to the normal equation

$$\dot{\mathbf{A}}(\eta_1,\ldots,\eta_k) = \mathbf{T}$$

- The coordinate ascent method is the bisection algorithm applied to each coordinate in the iterative recycling manner.
 - Start with an arbitrary (
 ⁿ₁,...,
 ⁿ_k) (some at hoc estimation, method
 of moments or plug-in estimation are recommended)
 - Solve the equation for the first derivative with respect to η₁ leading, say, to η₁¹ and replace η₁⁰ by it.
 - Repeat the previous step for the first derivative with respect to η₂ and so on until all η_i⁰'s are replaced by the corresponding η_i¹'s.
 - Repeat the above series of steps with $(\hat{\eta}_1^{j-1}, \ldots, \hat{\eta}_k^{j-1})$ replaced by $(\hat{\eta}_1^j, \ldots, \hat{\eta}_k^j)$ and leading to $(\hat{\eta}_1^{j+1}, \ldots, \hat{\eta}_k^{j+1})$ until the desired accuracy is attained.

Why does it work? - mathematics

Theorem 2.4.2.

If the conditions of Theorem 2.3.1 hold, then $(\hat{\eta}_1^r, \ldots, \hat{\eta}_k^r) \stackrel{r \to \infty}{\to} (\hat{\eta}_1, \ldots, \hat{\eta}_k)$.

A sketch of the proof.

- The log-likelihood is $\mathbf{t}_0^\top \eta A(\eta) + \log h(\mathbf{x})$ and it is a concave function.
- The algorithm finds the maximum along one coordinates while all other are fixed so that the value of the likelihood is increasing at each step of the algorithm.
- The log-likelihood is bounded in the interior of the set of the parameters thus there has to be a limit *l*(*η̂^r*) → λ ∈ (−∞, ∞)
- This implies that the
 ^γ for large r must reside in some closed ball in the interior
 of the parameter space and thus needs to have a convergent subsequence.
- The limit of this subsequence must be the same due to the strict concavity and the continuity of the derivatives of *A*, that take the value zero at each 'loop' of the algorithm so in the limit they have to be zero at all the coordinates.

Why does it work? – a picture



The coordinate ascent algorithm. The graph shows log likelihood contours, that is values of (θ_1, θ_2) where the loglikelihood is constant. At each stage with one coordinate fixed, find that member of the family of contours to which the vertical (or horizontal) line is tangent.

The Newton-Raphson Algorithm

Let us start with a visualization of the algorithm



The algorithm that, in general, can be shown to be faster than coordinate ascent. This method requires computation of the inverse of the Hessian, which may counterbalance its advantage in speed of convergence when it does converge.

NRA for MLE and exponential family

• For a general log-likelihood

$$\hat{\theta}_{new} = \hat{\theta}_{old} - \ddot{I}^{-1} (\hat{\theta}_{old}) \dot{I} (\hat{\theta}_{old}).$$

For the canonical exponential family

$$\hat{\eta}_{\textit{new}} = \hat{\eta}_{\textit{old}} - \ddot{\textit{A}}^{-1}(\hat{\eta}_{\textit{old}}) \left(\dot{\textit{A}}(\hat{\eta}_{\textit{old}}) - \mathbf{t}_0\right).$$

When likelihoods are non-concave, methods such as bisection, coordinate ascent, and Newton–Raphson's are still employed, though there is a distinct possibility of nonconvergence or convergence to a local rather than global maximum.

The EM algorithm – foundation

There are many models that have the following structure.

- There are ideal complete observations, X with density p(x, θ) with log likelihood l_x(θ) is easy to maximize.
- Unfortunately, we observe incomplete data given by S = S(X) with density q(s, θ), where the log-likelihood l_s(θ) is difficult to maximize.
- A fruitful way of thinking of such problems is in terms of S being incomplete data and representing a part of X, the rest of X is "missing" and its "reconstruction" is part of the process of estimating θ by maximum likelihood.
- The algorithm was formalized with many examples in Dempster, Laird, and Rubin (1977), though an earlier general form goes back to Baum, Petrie, Soules, and Weiss (1970).

Example – normal distribution with missing data

- Let $(Z_1, Y_1), \ldots, (Z_n, Y_n)$ be i.i.d. as $N(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2, \rho)$.
- Suppose that some of the *Z_i* and some of the *Y_i* are missing as follows:
 - For $1 \le i \le n_1$ we observe both Z_i and Y_i ,for $n_1 + 1 \le i \le n_2$, we observe only Z_i , and for $n_2 + 1 \le i \le n$, we observe only Y_i . The observed data are denoted by S.
 - In this case a set of sufficient statistics for the complete data is

$$T = (\overline{Z}, \overline{Y}, \overline{Z^2}, \overline{Y^2}, \overline{ZY}).$$

• One wants to reconstruct the missing parts needed for *T* from the incomplete data given in *S*.

Formulation of the EM algorithm

Let

$$J(heta| heta_0) = E_{ heta_0} \left(\log rac{oldsymbol{p}(X, heta)}{oldsymbol{p}(X, heta_0)} \Big| oldsymbol{S} = oldsymbol{s}
ight).$$

- **E-step** evaluate $J(\theta|\theta_0)$ as a function of θ .
- M-step maximize $J(\theta|\theta_0)$ over θ to get a 'new' θ_0 .
- repeat the above until the convergence.

If θ_{new} and θ_{old} the values at the end and at the beginning of the loop in the algorithm, respectively. Then

 $q(s, \theta_{new}) \geq q(s, \theta_{old}).$

The formal argument is interesting (not very intuitive), see Lemma 2.4.1.

Algorithms

The EM algorithm for exponential families

Theorem 2.4.3

For a canonical exponential family generated by (T, h) satisfying the conditions of Theorem 2.3.1. Let S(X) be any statistic, then the EM algorithm consists of the alternation

$$\hat{A}(\theta_{new}) = E_{\theta_{old}}(T(X)|S(X) = s).$$

Example (cont.): We can see that

$$\dot{A}(\theta) = E_{\theta}T = (\mu_1, \mu_2, \sigma_1^2 + \mu_1^2, \sigma_2^2 + \mu_2^2, \sigma_1\sigma_2\rho + \mu_1\mu_1)$$

so the left-hand side is straightforward. The right hand side can be easily derived by noting the well-known relations

$$E_{\theta}(Y|Z) = \mu_2 + \rho \sigma_2 (Z - \mu_1) / \sigma_1$$

$$E_{\theta}(Y^2|Z) = (\mu_2 + \rho \sigma_2 (Z - \mu_1) / \sigma_1)^2 + (1 - \rho^2) \sigma_2^2$$

leading to the new paremater values being simply regular functions of the sufficient statistic T but evaluated at the expected values in the E-step of the algorithm. See Example 2.4.6 for details.