

MC Statistical Errors in ML Fits

This note presents an example of a binned maximum-likelihood (ML) fit where a data histogram is modeled with a combination of signal and background components. The signal component is specified by a parametric function, and the background contribution for each bin is estimated by Monte Carlo (MC). Because the MC sample has a finite number of events, the background estimates have statistical errors that should be incorporated into the fit. Here this is done by modeling the numbers of MC events found in each bin as Poisson values whose means are treated as additional nuisance parameters in the likelihood function.

Essentially the same problem is described by Barlow and Beeston [1], who treat the case where the signal and background distribution shapes are both estimated from Monte Carlo.

Suppose an experiment results in a histogram of a continuous variable x with N bins represented by a set of numbers $\mathbf{n} = (n_1, \dots, n_N)$. We can model these as Poisson variables with expectation values

$$E[n_i] = \nu_i . \tag{1}$$

Suppose that the signal process describes the distribution of x with a pdf $f(x; \boldsymbol{\theta})$ where $\boldsymbol{\theta}$ is a set of parameters whose values we wish to estimate. In general the measurement of x will be characterized by a response matrix that gives the probability to measure a value in bin i given that the true value was in event j ,

$$R_{ij} = P(\text{event measured in bin } i \mid \text{true value in bin } j) . \tag{2}$$

The efficiency is thus simply the probability to be measured in any bin given that the true value was in bin j :

$$\varepsilon_j = \sum_{i=1}^N R_{ij} . \tag{3}$$

For purposes of the example here we will assume that the response matrix is known.

The pdf $f(x; \boldsymbol{\theta})$ will predict the expected numbers of events in a histogram of the *true* value of x . This is represented with a set of numbers $\boldsymbol{\mu} = (\mu_1, \dots, \mu_M)$, whose values are not necessarily integers. The true histogram need not have the same number of bins as that of the observed value of x , although in the example here we take them to be equal, i.e., $M = N$. The expected number of events in the histogram of the observed value of x must take into account the smearing and limited efficiency represented by the response matrix R , and it must also include contributions from background events. We can write the expectation value of n_i , ν_i , as

$$\nu_i = \sum_{j=1}^M R_{ij} \mu_j + \beta_i . \tag{4}$$

Here β_i is the expected number of background events in bin i , and the value of μ_j is related to the pdf $f(x; \boldsymbol{\theta})$ by

$$\mu_j = \mu_{\text{tot}} \int_{\text{bin},j} f(x; \boldsymbol{\theta}) dx . \quad (5)$$

Here μ_{tot} is the expected total true number of events, which is treated together with $\boldsymbol{\theta}$ as an unknown parameter.

In this example we will suppose that we have an MC model for the background. By subjecting the generated background events to the same detector simulation and selection criteria as used for the real data, we find a number of events m_i in the i th bin of x . Suppose now that the effective luminosity for the background sample is related to that of the real data (both assumed known) by

$$\tau = \frac{L_{\text{MC}}}{L_{\text{data}}} . \quad (6)$$

In principle we can make τ arbitrarily large by generating a larger MC sample, but in practice its value is often of order unity and it is rarely greater than 10. The number of events m_i in bin i of the background histogram can be modeled as a Poisson variable with expectation value

$$E[m_i] = \tau \beta_i . \quad (7)$$

Strictly speaking we should model m_i as following a binomial distribution with success probability p_i out of N_{MC} generated events, with $E[m_i] = p_i N_{\text{MC}}$. Here we can assume, however, that the analysis has been designed to suppress background events so that all of the p_i are small for the observed bins. We can therefore accurately model the m_i as Poisson variables.

Using a Poisson model for both the data n_i and background MC values m_i , we can write the likelihood function as

$$L(\mu_{\text{tot}}, \boldsymbol{\theta}, \boldsymbol{\beta}) = \prod_{i=1}^N \frac{\nu_i^{n_i}}{n_i!} e^{-\nu_i} \prod_{i=1}^N \frac{(\tau \beta_i)^{m_i}}{m_i!} e^{-\tau \beta_i} . \quad (8)$$

The log-likelihood is thus

$$\ln L(\mu_{\text{tot}}, \boldsymbol{\theta}, \boldsymbol{\beta}) = \sum_{i=1}^N (n_i \ln \nu_i - \nu_i + m_i \ln(\tau \beta_i) - \tau \beta_i) + C , \quad (9)$$

where C is a constant not depending on the parameters that can be dropped. The parameters μ_{tot} and $\boldsymbol{\theta}$ enter through the ν_i according to Eqs. (4) and (5), and the background values β_i enter directly through Eq. (4).

In general one maximizes the log-likelihood numerically. The difficulty one encounters here is that all N of the background values β_i all count as adjustable parameters. If one generalizes to the case where there are several background components, then there will be N additional parameters for each component. So a problem that initially had only a few parameters (the components of $\boldsymbol{\theta}$), has become one with potentially many tens or even hundreds of parameters.

In fact this is not as difficult as it may seem, since one has an excellent first guess for the value of β_i , namely, m_i/τ . And by using the method of Barlow and Beeston [1], the

numerical maximization of $\ln L$ is feasible even for cases with a very large number of bins times background components. In applying their method to this problem, one simply sets the derivatives of $\ln L$ with respect to $\boldsymbol{\theta}$ and $\boldsymbol{\beta}$ equal to zero to find the maximum, i.e.,

$$\frac{\partial \ln L}{\partial \theta_i} = \sum_{j=1}^N \left(\frac{n_j}{\nu_j} - 1 \right) \frac{\partial \nu_j}{\partial \theta_i} = 0, \quad (10)$$

$$\frac{\partial \ln L}{\partial \mu_{\text{tot}}} = \sum_{i=1}^N \sum_{j=1}^M \left(\frac{n_i}{\nu_i} - 1 \right) R_{ij} = 0, \quad (11)$$

$$\frac{\partial \ln L}{\partial \beta_i} = \frac{n_i}{\nu_i} - 1 + \frac{m_i}{\beta_i} - \tau = 0. \quad (12)$$

In practice, Eqs. (10) and (11) are not needed as the maximization with respect to these parameters is done numerically, e.g., with the program MINUIT. Solving Eq. (12) for β_i gives

$$\beta_i = \frac{m_i}{\tau + 1 - n_i/\nu_i}. \quad (13)$$

The strategy for maximizing $\ln L$ is to begin with the estimate for β_i that would be obtained for $\nu_i = n_i$, namely,

$$\beta_i = \frac{m_i}{\tau}, \quad (14)$$

and with these values fixed, to maximize $\ln L$ with respect to the parameters $\boldsymbol{\theta}$ and μ_{tot} . Using Eq. (4) with the updated parameters, one finds new values for the ν_i . If these ν_i represented the solution, then

$$\delta_i = \frac{m_i}{\tau + 1 - n_i/\nu_i} - \beta_i \quad (15)$$

would be zero. One can then update the estimates of the β_i using

$$\beta_i \rightarrow \beta_i + \eta \delta_i, \quad (16)$$

where η is a “learning rate” that can be adjusted to improve the convergence. Using $\eta = 1$ corresponds to taking Eq. (13) itself to give the updated values of β_i . In the example shown below, however, this led to an oscillating solution. Reasonable convergence was found after several iterations with $\eta = 0.2$.

In principle one can iterate as described above until the solution converges with the desired accuracy. At each iteration, the numerical maximization only involves the parameters $\boldsymbol{\theta}$ and μ_{tot} . In practice it was quicker to terminate the procedure after several iterations and then to use this trial solution in a numerical maximization of all of the parameters: $\boldsymbol{\theta}$, μ_{tot} and $\boldsymbol{\beta}$. The optimal trade-off between iterative and full numerical maximization will depend in general on the problem, and in particular on the number of components in $\boldsymbol{\theta}$ and $\boldsymbol{\beta}$.

The method above was applied to the Monte Carlo example shown in Fig. 1. The pdf $f(x)$ for the signal was a Gaussian with a mean of 5.0 and standard deviation of 1.0, truncated between the limits $0 \leq x \leq 10$. The parameter μ_{tot} representing the expected number of signal

events was 500. The background followed an exponential with mean 2.0 truncated in the same range. The expectation for the total number of background events was $\beta_{\text{tot}} = \sum_i \beta_i = 1000$. The solid histogram in Fig. 1 represents the generated data. for the fit in Fig. 1, the value of the parameter τ was taken to be equal to 1.0.




Figure 1: Result of the maximum-likelihood fit using $\tau = 1$. The solid histogram represents the data, the dotted curve is the fitted signal, the lower dotted histogram is the estimated background, and the upper dotted histogram is the sum of fitted signal and background (see text).

For the fit in Fig. 1 it was in fact possible to do the full numerical maximization of all parameters, including background values, when starting from the initial estimates $\beta_i = m_i/\tau$. To ensure convergence it is important to place limits on the β_i so that they remain non-negative, and convergence is improved if they are given a small positive value (e.g., 0.1) in the bins where $m_i = 0$.

The iterative procedure described above was effective, however, and reasonable convergence was found after 4 iterations using $\eta = 0.2$ for the learning rate. Taking $\eta = 1.0$ led to an oscillating solution for bins with large numbers of entries. At the end of the iterative phase, the full numerical maximization was carried out. Initial, intermediate and final fit results for the background parameters are shown in Fig. 2.




Figure 2: Result of the fit of the background parameters β . The circles show the initial estimates $\beta_i = m_i/\tau$ (with $\tau = 1$), the squares indicate the estimates after four iterations with the procedure described above, and the triangles are the final estimates after a full numerical maximization of $\ln L$ with respect to all of the parameters.

Figures 3 show the results when taking (a) $\tau = 0.1$ and (b) $\tau = 1.0$. The former represents a case where there are very large statistical uncertainties in the background estimates, and




Figure 3: Fit results using (a) $\tau = 0.1$ and (b) $\tau = 10$.

this is reflected in the full result (upper dotted histogram) following very closely the data histogram. That is, in the case $\tau = 0.1$, most of the information about the background is coming from the data histogram and not from the subsidiary background measurements. This results in larger statistical uncertainties for the parameters of interest. These are shown in Table 1 $\tau = 0.1, 1$ and 10 . In the case $\tau = 10$, the subsidiary measurements m_1, \dots, m_N determine the background very accurately, and this is reflected in the smaller statistical errors in the parameters of interest.

Table 1: Fitted values for the mean and width of the signal pdf and the parameter μ_{tot} for different values of the scale parameter τ .

τ	mean	width	μ_{tot}
0.1	5.275 ± 0.129	0.684 ± 0.147	399.1 ± 54.2
1.0	5.080 ± 0.082	0.979 ± 0.097	494.5 ± 33.8
10	5.147 ± 0.076	0.956 ± 0.079	469.2 ± 28.3

References

- [1] Roger Barlow and Christine Beeston, *Fitting using finite Monte Carlo samples*, Comp. Phys. Comm. 77 (1993) 219–228.