

On Bayesian Treatment of Systematic Uncertainties in Confidence Interval Calculation

Fredrik Tegenfeldt ^{a,b} and Jan Conrad ^{b,*}

^a*Iowa State University, Ames, IA 5011-3160, USA*

^b*CERN, PH-EP Department, CH-1211 Geneva 23*

Abstract

In high energy physics, a widely used method to treat systematic uncertainties in confidence interval calculations is based on combining a frequentist construction of confidence belts with a Bayesian treatment of systematic uncertainties. In this note we present a study of the coverage of this method for the standard Likelihood Ratio (aka Feldman & Cousins) construction for a Poisson process with known background and Gaussian or log-Normal distributed uncertainties in the background or signal efficiency. For uncertainties in the signal efficiency of upto 40 % we find over-coverage on the level of 2 to 4 % depending on the size of uncertainties and the region in signal space. Uncertainties in the background generally have smaller effect on the coverage. A considerable smoothing of the coverage curves is observed. A software package is presented which allows fast calculation of the confidence intervals for a variety of assumptions on shape and size of systematic uncertainties for different nuisance parameters. The calculation speed allows experimenters to test the coverage for their specific conditions.

Key words: Confidence Intervals, Systematic Uncertainties, Frequentist Methods,

1 Introduction

The calculation of confidence intervals in case of presence of systematic uncertainties is an open problem. Systematic uncertainties are those uncertainties present in parameters which will affect the calculated confidence interval, but which are not of prime interest, so called *nuisance parameters*. Examples for nuisance parameters are the efficiency or the predicted background rate.

In 1992 Cousins & Highland [1] proposed a method which is based on a Bayesian treatment of the nuisance parameters. They proposed to perform a frequentist construction (following Neyman [2]) and to replace the probability density function (PDF) describing the statistical process of prime interest with a PDF which is obtained by a convolution the original PDF with the one describing the uncertainties in the nuisance parameters:

$$P(n|s) \longrightarrow \int P(n|s, \epsilon') P(\epsilon'|\epsilon) d\epsilon' \quad (1)$$

where ϵ' is the true value of the nuisance parameter, ϵ denotes its estimate and s and n symbolize the signal hypothesis and the experimental outcome respectively.

Highland & Cousins only treated the case of Gaussian uncertainties in the signal efficiency. The method has since been generalized to operate with the

* Corresponding author: Jan Conrad, PH-EP Dept, F01910, CH-1211 Geneva 23
Email addresses: `Fredrik.Tegenfeldt@cern.ch` (Fredrik Tegenfeldt),

`Jan.Conrad@cern.ch` (Jan Conrad).

modern unified ordering scheme proposed by Feldman & Cousins [3] and taking into account several nuisance parameters (efficiencies and background) and correlations [5]. This generalized method has already been used in a number of particle and astroparticle physics experiments (e.g. [6] [7] [8], [9], [10]).

The most crucial property of methods for confidence interval construction is the coverage, which states:

A method is said to have coverage $(1-\alpha)$ if, in infinitely many repeated experiments the resulting confidence interval includes the true value with probability $(1-\alpha)$ irrespective of what the true value is.

100(1- α)% is hereby commonly taken to be 68% 90%, 95% etc. Recently, the coverage properties of a fully frequentist method, the Profile Likelihood method, have been studied [11]. The Profile Likelihood method was found to have surprisingly good coverage properties with a small (mostly negligible) amount of undercoverage.

In this note we undertake a systematic study of the coverage of the Bayesian method in order to give more substance to further recommendations on what method to use to calculate confidence intervals in presence of systematic uncertainties. Previous studies [12] of this method dealt only with certain limiting cases and were constrained by computational requirements.

The note is organized as follows: in the next section we will review the frequentist construction of confidence intervals, in particular the likelihood ratio ordering scheme. We will describe the Bayesian method to incorporate systematics in section 3. The C++ library used to perform the calculations is described in section 4. Coverage tests and connected subtleties are discussed in section 5. The final section is devoted to discussion and conclusion.

2 Frequentist confidence interval construction

Let us consider a Poisson probability density function (PDF), $p(n)_{s+b}$, for a fixed but unknown signal, s , in the presence of a known background with mean b . For every value of s we can find two values n_1 and n_2 such that

$$\sum_{n'=n_1}^{n_2} p(n')_{s+b} = 1 - \alpha \quad (2)$$

where $1 - \alpha$ denotes the confidence level (usually quoted as a $100(1-\alpha)\%$ confidence interval). Since we assume a Poisson distribution, the equality will generally not be fulfilled exactly. A set of intervals $[n_1(s+b, \alpha), n_2(s+b, \alpha)]$ is called a *confidence belt*. Graphically, upon a measurement, n_o , the *confidence interval* $[s_1, s_2]$ is determined by the intersection of the vertical line drawn from the measured value n_o and the boundary of the confidence belt. The probability that the confidence interval will contain the true value s is $1 - \alpha$, since this is true for all s per construction. The choice of the n_1 and n_2 is, however, not unique to define the confidence belt. An additional criterion has to be applied. The currently recommended ordering scheme [3] [4] of the elements in the sum in equation (2) is based on likelihood ratios. This approach automatically provides central confidence intervals when motivated and upper limits when necessary, therefore it is often denoted as the “unified approach”. The following algorithm is applied in solving equation (2):

For each n the s_{best} is found which maximizes the likelihood $\mathcal{L}(n)_{s+b}$. In case of a simple Poisson distribution with known background, s_{best} is given by $\max(0, n - b)$. Then for a fixed s the ratio

$$R(s, n)_{\mathcal{L}} = \frac{\mathcal{L}_{s+b}(n)}{\mathcal{L}_{s_{best}+b}(n)} \quad (3)$$

is computed for each n , and all n 's are consequently ranked according to the value of this ratio. Values of n are included in the confidence belt starting with the n with the highest rank (largest $R_{\mathcal{L}}$) and then decreasing rank until $\sum_{n=n_1}^{n_2} p(n)_{s+b} = 1 - \alpha$. After the confidence belt has been constructed in this way, the confidence interval $[s_1, s_2]$ is found as described above.

3 Incorporation of systematic uncertainties: the Bayesian way

The ordering schemes are unaffected by the way of treating systematic uncertainties considered here. As mentioned earlier, the PDF describing the statistical process will however be modified. Two concrete examples are the following: In the case that the only uncertainty present is a theoretical (assumed Gauss-shaped) uncertainty of the background process the PDF is modified to:

$$q(n)_{s+b} = \frac{1}{\sqrt{2\pi}\sigma_b} \int_0^{\infty} p(n)_{s+b'} e^{-\frac{(b-b')^2}{2\sigma_b^2}} db' \quad (4)$$

Here b is the estimated background level, and σ_b is the uncertainty in the background estimation. If, in addition to the theoretical uncertainty for background, there is the need to include the uncertainty in the signal detection efficiency the expression for $q(n)_{s+b}$ might be extended to:

$$q(n)_{s+b} = \frac{1}{2\pi\sigma_b\sigma_\epsilon} \times \int_0^{\infty} \int_0^{\infty} p(n)_{b'+\epsilon's} e^{-\frac{(b-b')^2}{2\sigma_b^2}} e^{-\frac{(1-\epsilon')^2}{2\sigma_\epsilon^2}} db' d\epsilon' \quad (5)$$

where σ_ϵ is the uncertainty in the detection efficiency expressed in *relative* terms with respect to the nominal efficiency¹. It is important to realize that

¹ of course the efficiency could be defined in absolute terms as well.

the integration variables, here ϵ' and b' , are the possible “true” (but unknown) values of nuisance parameter. This indicates that this method is based on Bayesian statistics.

Some examples for confidence intervals computed by this method are shown in table 1.

4 pole++

For the coverage studies presented in this paper a reasonably fast and efficient code is required. Hence, a user-friendly and flexible C++ library of classes was developed based on the FORTRAN routine presented in [13]. The library is independent of external libraries and consists of two main classes, *Pole* and *Coverage*. The first class takes as input the number of observed events, the efficiency and background with uncertainties and calculates the limits using the method described in this paper. The integrals are solved analytically. *Coverage* generates user-defined pseudoexperiments and calculates the coverage using *Pole*. Presently the library supports Gauss and log-Normal distributed PDF for description of the nuisance parameters. Flat and user-defined PDFs are about to be implemented as well as correlations for the Gauss case. The class is dynamically optimized depending on if one wishes to calculate single (or few) confidence intervals or if one wants to perform a coverage study. Without these optimisations the calculation of a single interval takes about 1 second (wall clock time) on a 1 GHz Pentium III processor. The duration of a full coverage study (typically requiring the calculation of $\mathcal{O}(10^5)$ confidence intervals) ranges between a couple of minutes for small uncertainties and small signal hypotheses to order of 10 hours for large uncertainties and high signal

hypotheses. The performance of a coverage study thus seems feasible for the particular set of systematic uncertainties that may appear in real experiments. The pole++ library can be obtained from <http://cern.ch/tegen/statistics.html>

5 Coverage Studies

The coverage of the method is studied using MC simulations of pseudo-experiments with given true value of the prime parameter (signal) or nuisance parameter (efficiency, background). The *estimated* values of the nuisance parameters were assumed to be Gaussian or log-Normal distributed around the given true value. The outcome of one experiment thus consisted of a number of observed events (following a Poisson distribution with known background and depending on the true efficiency and true background) and the estimate of the nuisance parameter(s).

Figure 1 shows the coverage² as a function of signal hypothesis for two different sizes of uncertainties (5 % and 40 %). The uncertainty considered is in the signal efficiency and assumed to be Gauss-shaped. It can be seen that the Bayesian method causes over-coverage which is larger for larger systematic uncertainty. This is also reflected in figure 3 where the mean coverage (mean taken over all tested signal hypotheses) is shown as a function of assumed uncertainty. For Gaussian uncertainties in the signal efficiency we find an increase in mean coverage by $\sim 1\%$ in the uncertainty range between 10 % and 40 %.

The see-saw structure which is generally seen in the no-uncertainty case (due

² we will denote the coverage calculated by the MC simulations as $(1 - \alpha)_{eff}$ to distinguish it from the nominal coverage $(1 - \alpha)$

to the discrete experimental outcome) is considerably smoothed when systematic uncertainties are introduced. This is exemplified in figure 2, which compares the coverage curve with 5 % uncertainties on the signal efficiency with the zero uncertainties case. To quantify this further we present the rms of the coverage together with the mean in figure 3. The smoothing is due to the fact that we add a continuous variable to the problem, meaning there are more degrees of freedom to fulfil the sum condition of equation 2. The effect of the smoothing is that whereas for some signal hypothesis the coverage is increased for others it is decreased with respect to the zero uncertainty case. The mean coverage is therefore only rather weakly dependent on the systematic uncertainties. For higher signal hypotheses, where the effect of the smoothing is less pronounced, the effect of the uncertainties is therefore stronger. Taking the mean over signal hypotheses $s_{true} > 6$ the mean coverage increases from 92 % (at zero uncertainties) to 94 % (at 40 % uncertainties). A side effect is that the introduction of rather small uncertainties seems to improve the coverage with respect to the zero uncertainty case for parts of the tested hypothesis space (see e.g. figure 2).

Figure 4 shows the coverage for two different sizes of Gaussian uncertainty on the background estimate. Except for the smoothing effect the coverage is seemingly independent of the size of the uncertainties. An uncertainty in the background will yield similar results to an uncertainty in the signal efficiency only in the regime where the signal hypotheses are of similar size as the background expectation. For larger signal hypotheses the coverage curve for the uncertainty in the background will be approaching the zero uncertainty case, which is reflected in a slight slope of the coverage curve for high uncertainty in figure 4. The result is thus that a mean coverage depending on uncertainties in the background will be less affected than the corresponding curve with

uncertainties in the signal efficiency. We show the mean coverage as well as the rms in figure 5.

In figure 6 the effect of having to consider uncertainties both in signal and background detection efficiency is visible. At large signal hypotheses the coverage plot is dominated by the uncertainty in signal efficiency, at low signal hypotheses coverage benefits from the smoothing effect, since we added yet another degree of freedom.

It should be noted that for large uncertainties the Gaussian model is not appropriate. The reason is that a Gaussian model will then significantly extend to the unphysical region of the space of experimental outcomes. In the Bayesian treatment this case is dealt with by truncating the Gauss distribution at zero³, and this is consequently the way it is dealt with during the coverage test. Considering the coverage test there is an additional subtlety: in order for the measured efficiency to be a maximum likelihood estimate of the true efficiency we would have to renormalize the Gauss distribution to the truncated Gauss distribution in order to obtain the correct PDF. However, instead of doing this a posteriori fix, it is more reasonable to use a log-Normal distribution to model the uncertainty in the efficiency. The coverage for the log-Normal distribution is shown in figure 7. As can be seen for the highest uncertainties considered in this note, the Gauss distribution is still a very good approximation to the log Normal model.

³ it is worth noting that Cousins & Highland develop their method for unbiased estimators, i.e. the problem would not be treated in this way.

6 Discussion and Conclusions

In this note we presented coverage studies for the Bayesian treatment of systematic uncertainties. One overall conclusion is that the Bayesian treatment leads some over-coverage. However, introducing a continuous nuisance parameter into the discrete Poisson problem results in a smoothing of the coverage curves. The mean coverage is therefore only weakly affected by the Bayesian treatment of nuisance parameters and under certain circumstances even improves with respect to the zero uncertainty case.

In a frequentist approach the measured estimate of a nuisance parameter is considered to be distributed around a given true value, which is consequently the way coverage was calculated. The Bayesian method on the other hand views the true value as distributed around the measured value. The underlying assumption for going from one approach to the other is (at least in case of a Gauss-distribution) a flat prior probability of hypotheses. The present study indicates that this assumption does not lead to a violation (except for over-coverage) of the coverage requirement.

The routines used for the presented calculations are reasonably fast and publicly available. A coverage study is therefore feasible for each problem at hand. The confidence level required for the confidence intervals could then in principle be adjusted to recover correct coverage.

7 Acknowledgements

The authors thank Robert Cousins for a very useful discussion.

References

- [1] R. D. Cousins and V. L. Highland, Nucl. Instrum. Meth. **A320**, 331, (1992).
- [2] J. Neyman, Phil. Trans. Royal Soc. London **A**, 333, (1937).
- [3] G. J. Feldman and R. D. Cousins, Phys. Rev **D57**, 3873, (1998).
- [4] A. Stuart and J. K. Ord: Kendall's Advanced Theory of Statistics, Vol. 2, Classical Inference and Relationship, Oxford University Press, New York (1991).
- [5] J. Conrad, O. Botner, A. Hallgren and C. Perez de los Heros, Phys. Rev. D **67** (2003) 012002 [arXiv:hep-ex/0202013].
- [6] B. Abbott *et al.* [LIGO Collaboration], Phys. Rev. D **69** (2004) 102001.
- [7] Y. Chao *et al.* [Belle Collaboration], Phys. Rev. D **69** (2004) 111102 [arXiv:hep-ex/0311061].
- [8] K. Eguchi *et al.* [KamLAND Collaboration], Phys. Rev. Lett. **92** (2004) 071301 [arXiv:hep-ex/0310047].
- [9] I. Abt *et al.* [HERA-B Collaboration], arXiv:hep-ex/0405059.
- [10] J. Ahrens [AMANDA Collaboration], Phys. Rev. Lett. **92** (2004) 071102 [arXiv:astro-ph/0309585].
- [11] W. A. Rolke, A. M. Lopez, J. Conrad and F. James arXiv:physics/0403059. 098301, (2000).
- [12] J. Conrad, O. Botner, A. Hallgren and C. Perez de los Heros, published in Proc. of Conference on Advanced Statistical Techniques in Particle Physics, Durham, England, March 2002
- [13] J. Conrad, Computer Physics Communications **158** 117-123 (2004)

8 Figures and tables

n_0	b	rel. signal efficiency uncertainty	Likelihood Ratio interval
2	2	0	0: 3.91
		0.2	0: 3.89
		0.4	0: 4.68
4	2	0	0: 6.59
		0.2	0: 7.16
		0.4	0: 8.99
6	2	0	1.08: 9.47
		0.2	1.07:10.09
		0.4	1.02:13.31

Table 1

Examples of likelihood ratio 90% confidence intervals with Bayesian treatment of systematic uncertainties. Uncertainties are assumed to be Gauss distributed in the signal efficiency.




Fig. 1. Calculated coverage as function of signal hypothesis. Two case are shown: 5 % and 40 % Gaussian uncertainties in the signal efficiency. The nominal coverage was 90%.




Fig. 2. Calculated coverage with 5 % systematic uncertainties is compared with the zero uncertainties case. The nominal coverage was 90 %.




Fig. 3. Calculated mean coverage and rms as function of Gaussian shape uncertainties in the signal efficiency. The mean was calculated over the full range of signal hypotheses and for signal hypotheses larger than 6, respectively. The rms was calculated over the full range of signal hypotheses. The nominal coverage was 90 %.




Fig. 4. Calculated coverage as function of signal hypothesis. Two case are shown: 5 % and 40 % Gaussian uncertainties in the background prediction. The nominal coverage was 90 %.

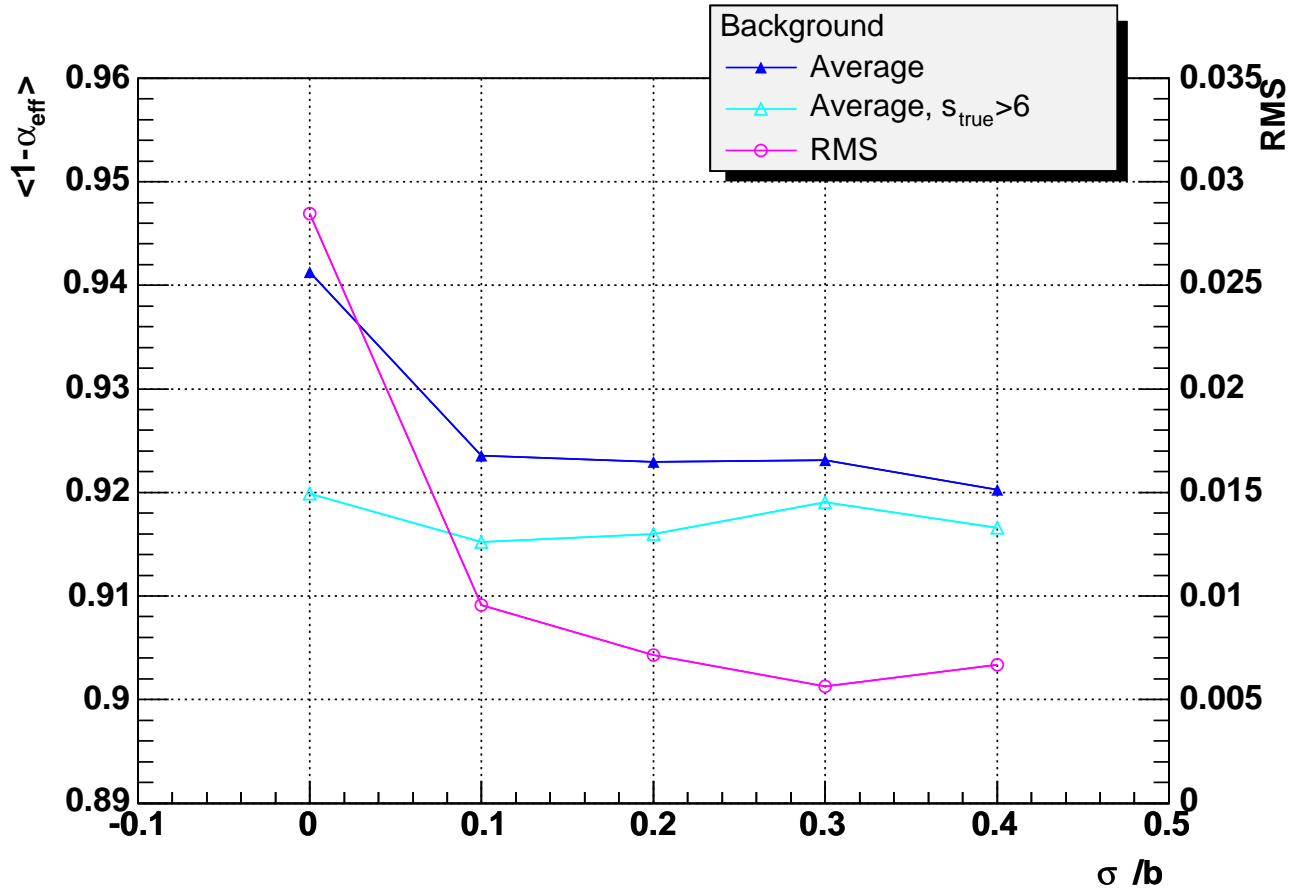


Fig. 5. Calculated mean coverage and rms as function of Gaussian shape uncertainties in the background prediction. The mean was calculated over the full range of signal hypotheses and for signal hypotheses larger than 6, respectively. The rms was calculated over the full range of signal hypotheses. The nominal coverage was 90 %

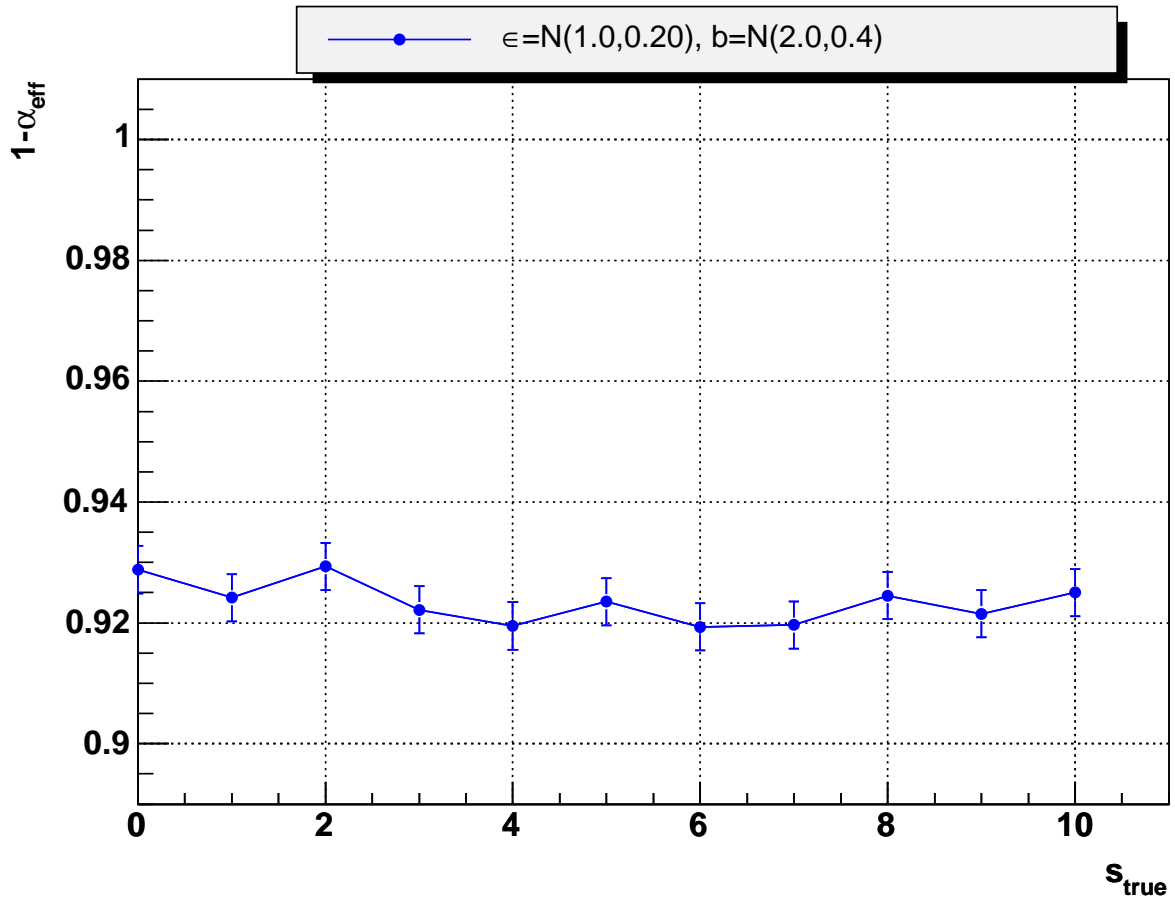


Fig. 6. Coverage for the case that uncertainties are present in both the signal detection efficiency and in the background prediction. The nominal coverage was 90 %




Fig. 7. Coverage as a function of signal hypothesis for a Gaussian PDF and a log Normal PDF