/CDF/PUB/EXOTIC/PUBLIC/4476

Poisson Upper Limits Incorporating Uncertainties in Acceptance and Background

John Conway
Rutgers University

Kaori Maeshima
Fermi National Accelerator Laboratory

8 February 1998

Abstract

This note presents a technique for the calculation of Poisson upper limits at some desired confidence level given a certain number of observed events, an expected number of background events, and uncertainty on both the overall acceptance and the expected background. This hybrid frequentist/Bayes technique represents the standard adopted by the CDF Collaboration in presenting the results of new particle search counting experiments. A “fully Bayesian” approach incoporating uncertainties yields limits which match the frequentist limits in the zero-uncertainty case, but otherwise exceeds the frequentist limits.

1 Motivation

In many, if not most, new particle searches in high energy physics, one selects from a large number of recorded events those which bear characteristics of the new process while minimizing the retention of events from well-understood processes. This typically results in a small number of events passing the selection requirements, consistent with the expectation from a calculation of the expected background. At this stage one can determine a Poisson upper limit on the number of signal events present in the sample, at some desired confidence level (usually 95%), employing the method described by the Particle Data Group (PDG). [1].

The PDG method, however, is silent on the issue of how to incorporate the uncertainties present in the expected number of background events and in the overall acceptance. Other authors [2, 3] have separately discussed methods for incorporating uncertainties in the background or acceptance, but none to date (to our knowledge) have presented a method incorporating both.

This note discusses the method used by the CDF Collaboration to determine Poisson upper limits in the presence of uncertainties (both statistical and systematic) simultaneously in the accpetance and background, essentially merging the approaches of Zech and Huber *et al.* referred to above. Though often thought of as an example of a classical (“frequentist”) statistical procedure, this method in fact uses Bayesian notions of a probability density function (pdf) for the unknown true values of the acceptance and expected background. In the case of no uncertainties in the acceptance or expected background, the full-blown Bayesian treatment gives the same upper limits as the frequentist PDG method.[3] As shown below, though, when uncertainties are incorporated into the fully Bayesian algorithm, the resulting limits exceed those obtained in the frequentist method.

This note should serve as a reference for this method for those who wish to cross check their own calculations or wish to understand and use the computer programs discussed herein.

2 Upper Limits Without Uncertainties

Given n_0, an (integer) number of observed events, the probability P for observing *that number* depends on μ, the mean number of events expected, according to the Poisson distribution:

$$P(n_0;\mu) = \frac{\mu^{n_0} e^{-\mu}}{n_0!} \quad . \tag{1}$$

In new particle searches one wishes to determine the value of μ. We define the Poisson upper limit N on the number of expected events[1] as that value of μ for which there is some probability ϵ to observe n_0 or fewer events. The confidence level (CL) of the upper limit is then simply $1-\epsilon$. One can calculate ϵ by summing over the Poisson probabilities:

$$\epsilon = \sum_{n=0}^{n_0} P(n;\mu) \quad . \tag{2}$$

In practice, then, to calculate N one varies μ until finding the value of ϵ corresponding to the desired CL; N is the resulting value of μ.

If one expects an average of μ_B background events among the n_0 observed, and if one knows μ_B precisely, then the PDG also discusses a method for calculating a Poisson upper limit N on the number of *signal* events present in the observation. The value of N represents that value of μ_S, the mean number of signal events expected, for which the probability is $1-\epsilon$ that in a random experiment one would observe *more than* n_0 events *and* have $n_B \leq n_0$, where n_B is the number of background events present in the sample. This can be calculated as before by adjusting N until the relation

$$\epsilon = \frac{\sum_{n=0}^{n_0} P(n;\mu_B+N)}{\sum_{n=0}^{n_0} P(n;\mu_B)} \tag{3}$$

obtains. As the PDG points out, this results in a "conservative" upper limit in that for some true μ_S the probability of obtaining $N > \mu_S$ exceeds $1-\epsilon$ on average. Also, if one obtains a value of n_0 significantly lower than μ_B, it is more conservative to assume $n_0 \approx \mu_B$ to obtain N.

[1]Note that N is a real number, not an integer; we follow the notation of the PDG here.

3 Incorporation of Uncertainties

Suppose that one knows the value of μ_B to within an overall (statistical plus systematic) gaussian uncertainty of σ_B, and the overall acceptance A to within an overall uncertainty of σ_A. In this case the relative uncertainty on μ_S is σ_A/A. One can define the Poisson upper limit N on μ_S as before: that value of the true μ_S for which one would observe *more than* n_0 events *and* have $n_B \leq n_0$. In this case, however, one seeks the value of N such that

$$\epsilon = \frac{\displaystyle\sum_{n=0}^{n_0} \frac{1}{\sqrt{2\pi\sigma_N^2}} \int_0^\infty \int_0^\infty P(n; \mu'_B + \mu'_S) e^{-\frac{(\mu_B - \mu'_B)^2}{2\sigma_B^2}} e^{-\frac{(N-\mu'_S)^2}{2\sigma_N^2}} d\mu'_B d\mu'_S}{\displaystyle\sum_{n=0}^{n_0} \int_0^\infty P(n; \mu_B) e^{-\frac{(\mu_B - \mu'_B)^2}{2\sigma_B^2}} d\mu'_B} \tag{4}$$

where we take $\sigma_N = N\sigma_A/A$. In this way one assumes an *a priori* gaussian distribution of the true values of μ_S and μ_B about the values obtained in subsidiary studies, with width given by the uncertainties obtained in those studies.

4 Frequentist Algorithm and Results

One can perform the formidable integral in equation 4 by various numerical techniques. In the method presented here, rather than performing the integral directly, for each test value of N one generates a large ensemble of random pseudo-experiments, varying the expected number of signal and background about their nominal values according to a gaussian distribution. In each experiment, the expected number of signal and background events are chosen from the gaussians, and Poisson-distributed numbers of signal (n_S) and background (n_B) events are generated. For those trials where $n_B \leq n_0$, the fraction f in which $n_B + n_S > n_0$ is recorded.

The confidence level for a given N is in fact equal to f; one must then simply vary N until the desired CL $(1-\epsilon)$ is obtained. This is implemented in the FORTRAN program `poilim.f`. [4] The program has two phases. In the first phase, the program makes an initial guess and begins generating pseudo-experiments, calculating f until either 1) the binomial error on f is such that the deviation is more than four standard deviations from the desired, or 2) the number of pseudo-experiments exceeds 10^6.

At this point the program makes a new guess, increasing or decreasing it by the present step size, and then decreasing the step size if the new guess is less than the previous one. Eventually, when the guess lies near the desired value, the loop will terminate because of the limit of 10^6 trials; at this point the program begins the second phase.

In the second phase the program scans the region near the latest guess, and then fits a quadratic to the values $f_i(N_i)$ obtained. From the fit, the value of $\tilde{N}$ where $f(\tilde{N}) = 1 - \epsilon$ can be calculated, including the uncertainty on $\tilde{N}$. Once this uncertainty is less than that specified by the user, the program terminates.

The accuracy of this method (as with any numerical estimate) is limited by compute time. The slope of $f(N)$ near the limit typically lies in the range 0.005-0.05 for 95% CL, and so to obtain an accuracy on N much better than 0.01 requires very long calculations. Also, the program makes use of the CERNLIB routine POISSN to generate Poisson-distributed random numbers. The amount of time this routine takes can vary by more than an order of magnitude depending on the case.

Tables 1-3 provide sample results, against which one can cross check other calculations. As the table shows, including the effect of the uncertainties makes the Poisson upper limits larger ("worse") than those in the zero-uncertainty case, as one would expect. Also, the results of this calculation agree with the exact results obtainable in the uncertainty-free case.

Note that Cousins [5] provides a simple formula for the case where the background is accurately known, and where one knows N_0, the upper limit in the uncertainty-free case. In the notation used here the formula can be written

$$N = N_0 \left(1 + \frac{1 - \left[(N_0 - n_0)^2 \sigma_A^2 / A^2 \right]^{1/2}}{N_0 - n_0} \right) \quad . \tag{5}$$

In those cases where the background is known precisely, use of this expression avoids the lengthy numerical calculation of `poilim`.

5 Bayesian Method and Results

One can also obtain upper limits on a Poisson process using a "fully Bayesian" approach, as discussed in the literature. [3] A Bayesian deems it sensible to treat the unknown expected number of signal events as a random variable, for which there is some "prior" probability density function (pdf) $\mathcal{P}(\mu_S)$. Given the observation of n_0 events, one can then construct a "posterior" pdf $\mathcal{P}(\mu_s; n_0)$ which depends on the likelihood $\mathcal{L}(n_0; \mu_s)$ for observing n_0 events given μ_S expected:

$$\mathcal{P}(\mu_S; n_0) = \frac{\mathcal{L}(n_0; \mu_S) \mathcal{P}(\mu_S)}{\int_0^\infty \mathcal{L}(n_0; \mu_S) \mathcal{P}(\mu_S) d\mu_S} \quad . \tag{6}$$

One can set a Bayesian upper limit (or any other confidence interval) on the unknown parameter μ_S, then, simply from integration of $\mathcal{P}(\mu_S; n_0)$.

The values obtained depend, of course, on the choice of the prior $\mathcal{P}(\mu_S)$. In considering the results of a particular experiment, usually one wishes to use an "uninformed" prior pdf; that is, one wants to give no *a priori* bias to certain values of μ_S. This usually results, then, in choosing $\mathcal{P}(\mu_S)$ to be uniform for all physical values of μ_S: $\mathcal{P}(\mu_S) =$const. for $\mu_S \geq 0$.[2]

Extension to the case where one expects μ_B background is straightforward:

$$\mathcal{P}(\mu_S; n_0, \mu_B) = \frac{\mathcal{L}(n_0; \mu_S + \mu_B) \mathcal{P}(\mu_S)}{\int_0^\infty \mathcal{L}(n_0; \mu_S + \mu_B) \mathcal{P}(\mu_S) d\mu_S} \quad . \tag{7}$$

For uniform prior $\mathcal{P}(\mu_S)$ this reduces to

$$\mathcal{P}(\mu_S; n_0, \mu_B) = \frac{\mathcal{L}(n_0; \mu_S + \mu_B)}{\int_0^\infty \mathcal{L}(n_0; \mu_S + \mu_B) d\mu_S} \quad . \tag{8}$$

Remarkably, as Cousins points out [3], the upper limits obtained with this expression match exactly those obtained with equation (3). Note also that the denominator of equation (7) can simply be regarded as a normalization constant whose value depends on n_0 and μ_B. Thus we see that

$$\mathcal{P}(\mu_S; n_0, \mu_B) \propto \mathcal{L}(n_0; \mu_S + \mu_B) \quad . \tag{9}$$

[2]Note that such a pdf is formally non-normalizable.

n_0	0	1	2	3	4	5
μ_B	0.000	0.000	0.000	0.000	0.000	0.000
σ_B	0.000	0.000	0.000	0.000	0.000	0.000
σ_A/A						
0.000	2.996	4.745	6.296	7.754	9.152	10.512
0.010	2.996	4.746	6.297	7.756	9.154	10.515
0.020	2.998	4.748	6.301	7.760	9.162	10.524
0.030	2.999	4.752	6.307	7.771	9.175	10.539
0.040	3.003	4.758	6.317	7.783	9.191	10.559
0.050	3.007	4.766	6.330	7.800	9.213	10.586
0.060	3.012	4.776	6.344	7.821	9.241	10.618
0.070	3.019	4.788	6.362	7.846	9.271	10.656
0.080	3.026	4.802	6.384	7.874	9.308	10.702
0.090	3.034	4.818	6.408	7.908	9.350	10.754
0.100	3.042	4.835	6.436	7.945	9.398	10.813
0.110	3.051	4.856	6.467	7.987	9.452	10.878
0.120	3.063	4.878	6.502	8.034	9.511	10.952
0.130	3.075	4.904	6.539	8.086	9.578	11.034
0.140	3.089	4.931	6.581	8.143	9.653	11.125
0.150	3.104	4.961	6.628	8.205	9.733	11.225
0.160	3.120	4.994	6.678	8.275	9.819	11.332
0.170	3.138	5.029	6.734	8.348	9.915	11.452
0.180	3.157	5.067	6.792	8.429	10.019	11.577
0.190	3.178	5.110	6.856	8.518	10.133	11.715
0.200	3.201	5.155	6.926	8.614	10.255	11.866
0.210	3.225	5.205	7.002	8.719	10.388	12.028
0.220	3.252	5.258	7.085	8.832	10.531	12.203
0.230	3.280	5.316	7.174	8.952	10.688	12.395
0.240	3.312	5.379	7.271	9.084	10.856	12.598
0.250	3.344	5.446	7.374	9.225	11.035	12.817

Table 1: "Frequentist" upper 95% C.L. limits (accurate to ±0.001) on a Poisson process with background, incorporating uncertainties by Bayesian integration.

n_0	1	1	1	1	1	1
μ_B	1.000	1.000	1.000	1.000	1.000	1.000
σ_B	0.000	0.100	0.200	0.300	0.400	0.500
σ_A/A						
0.000	4.112	4.116	4.129	4.150	4.173	4.190
0.010	4.113	4.117	4.130	4.150	4.173	4.191
0.020	4.115	4.120	4.132	4.152	4.176	4.192
0.030	4.119	4.123	4.136	4.156	4.180	4.197
0.040	4.123	4.127	4.140	4.161	4.185	4.202
0.050	4.130	4.133	4.146	4.167	4.191	4.208
0.060	4.137	4.141	4.154	4.175	4.198	4.216
0.070	4.146	4.150	4.164	4.184	4.208	4.226
0.080	4.157	4.161	4.173	4.194	4.219	4.236
0.090	4.170	4.174	4.186	4.207	4.231	4.248
0.100	4.183	4.188	4.200	4.221	4.245	4.263
0.110	4.198	4.203	4.216	4.236	4.261	4.278
0.120	4.216	4.220	4.233	4.254	4.278	4.297
0.130	4.234	4.239	4.251	4.272	4.298	4.316
0.140	4.254	4.259	4.272	4.294	4.318	4.337
0.150	4.277	4.282	4.294	4.316	4.341	4.360
0.160	4.302	4.306	4.320	4.341	4.366	4.385
0.170	4.328	4.334	4.347	4.370	4.393	4.414
0.180	4.358	4.363	4.376	4.398	4.424	4.444
0.190	4.389	4.394	4.408	4.432	4.457	4.476
0.200	4.424	4.428	4.443	4.465	4.491	4.511
0.210	4.461	4.465	4.480	4.503	4.531	4.550
0.220	4.500	4.505	4.520	4.543	4.571	4.593
0.230	4.544	4.549	4.564	4.588	4.615	4.636
0.240	4.592	4.597	4.611	4.636	4.664	4.685
0.250	4.643	4.647	4.662	4.687	4.716	4.738

Table 2: “Frequentist” upper 95% C.L. limits (accurate to ± 0.001) on a Poisson process with background, incorporating uncertainties by Bayesian integration.

n_0	1	2	3	4	5
μ_B	1.000	2.000	3.000	4.000	5.000
σ_B	0.300	0.300	0.300	0.300	0.300
σ_A/A					
0.000	4.150	4.856	5.424	5.915	6.354
0.010	4.150	4.857	5.424	5.916	6.355
0.020	4.152	4.859	5.427	5.919	6.358
0.030	4.156	4.863	5.433	5.924	6.363
0.040	4.161	4.869	5.439	5.931	6.372
0.050	4.167	4.877	5.447	5.941	6.382
0.060	4.175	4.886	5.458	5.952	6.396
0.070	4.184	4.897	5.471	5.966	6.411
0.080	4.194	4.910	5.487	5.982	6.428
0.090	4.207	4.925	5.504	6.002	6.448
0.100	4.221	4.941	5.523	6.023	6.470
0.110	4.236	4.959	5.543	6.046	6.495
0.120	4.254	4.980	5.567	6.072	6.524
0.130	4.272	5.003	5.593	6.101	6.555
0.140	4.294	5.029	5.621	6.133	6.588
0.150	4.316	5.058	5.652	6.167	6.624
0.160	4.341	5.088	5.686	6.205	6.666
0.170	4.370	5.122	5.723	6.245	6.711
0.180	4.398	5.157	5.763	6.288	6.759
0.190	4.432	5.195	5.808	6.337	6.811
0.200	4.465	5.236	5.856	6.390	6.867
0.210	4.503	5.281	5.908	6.447	6.928
0.220	4.543	5.330	5.964	6.508	6.994
0.230	4.588	5.382	6.024	6.573	7.064
0.240	4.636	5.440	6.088	6.644	7.140
0.250	4.687	5.502	6.155	6.722	7.222

Table 3: "Frequentist" upper 95% C.L. limits (accurate to ± 0.001) on a Poisson process with background, incorporating uncertainties by Bayesian integration.

number observed	relative uncertainty on acceptance	expected background	background uncertainty	classical upper limit	Bayesian upper limit
0	0.00	0.00	0.00	2.996	2.996
0	0.10	0.00	0.00	3.042	3.076
1	0.00	0.00	0.00	4.745	4.745
1	0.00	1.00	0.00	4.112	4.109
1	0.10	1.00	0.10	4.188	4.256
1	0.25	1.00	0.25	4.673	5.150
3	0.25	3.00	0.50	6.218	7.288
5	0.20	5.00	0.50	6.918	7.559

Table 4: Sample 95% CL "Frequentist" upper limits (accurate to ± 0.001) and Bayesian upper limits for comparison and cross checks.

To incorporate uncertainties on the signal and background one treats the expected background and signal as unknown parameters with uniform prior pdf, with Gaussian likelihood about the estimates from subsidiary studies, just as in the frequentist case. One thus obtains

$$\mathcal{P}(\mu_S; n_0, \mu_B) \propto \frac{1}{2\pi\sigma_B\sigma_S} \int_0^\infty \int_0^\infty \mathcal{L}(n_0; \mu'_S + \mu'_B) e^{-\frac{(\mu_B-\mu'_B)^2}{2\sigma_B^2}} e^{-\frac{(\mu_S-\mu'_S)^2}{2\sigma_S^2}} d\mu'_B d\mu'_S \quad , \tag{10}$$

where $\sigma_S = (\sigma_A/A)\mu_S$ comes from the relative uncertainty on the acceptance.

To calculate upper limits, one can simply calculate the right hand side of equation (9) for an appropriate range of μ_S, and then define the upper limit on μ_S as that value for which

$$\epsilon = \frac{\int_{\tilde{\mu_S}}^\infty \mathcal{P}(\mu_S; n_0, \mu_B) d\mu_S}{\int_0^\infty \mathcal{P}(\mu_S; n_0, \mu_B) d\mu_S} \tag{11}$$

obtains for some desired confidence level $1-\epsilon$. Such a calculation is implemented in the program `bayes.f`. [6]

Figure 1 shows the curve of $\mathcal{P}(\mu_S; n_0, \mu_B)$ obtained for the case $n_0 = 5$, $\mu_B = 4.5$, σ_B=0.5, and a relative uncertainty on the acceptance of 20%. Numerical integration of this curve indicates a 95% upper limit at 7.76 events. The classical method gives a value of 7.24 events; as Table I shows, in general the Bayes upper limits exceed those obtained classically, if the uncertainty on the background or acceptance is non-zero. Thus, applied to new particle searches, Bayesian limits will in general result in weaker exclusion. Table 4 illustrates this for a few sample points.

Acknowledgements

We wish to thank Peter Wilson for numerically cross checking the formula in equation 5 with the results of the program `poilim`.

References

[1] Particle Data Group, Phys. Rev. D 54, 164 (1996).

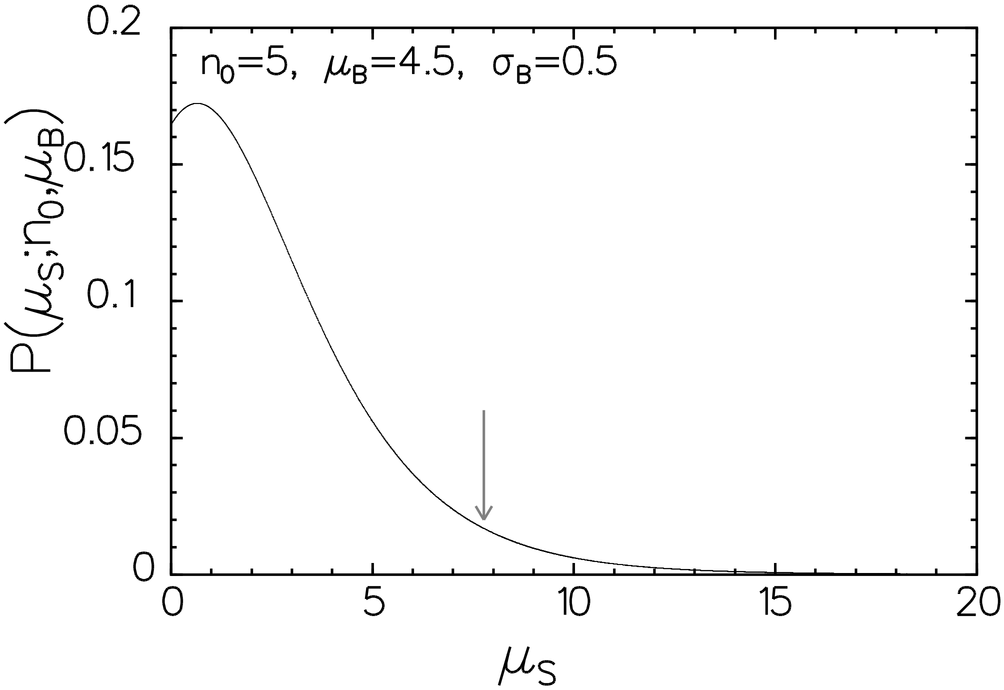

Figure 1: Plot of $\mathcal{P}(\mu_S; n_0, \mu_B)$ for the case $n_0 = 5$, $\mu_B = 4.5$, σ_B=0.5, and a relative uncertainty on the acceptance of 20%. The arrow indicates the point at 7.76, for which 5% of the area under the curve lies to the right.

[2] G. Zech, Nucl. Instrum. Methods Phys. Res., Sect. A 277, 608 (1989); T.M. Huber *et al.*, Phys. Rev. D 41, 2709 (1990).

[3] R. Cousins, Am. J. Phys. 63, 398 (1995) and references therein.

[4] See /home/lv59/conway/poilim.f.

[5] R. Cousins, V. Highland, Nucl. Instrum. Methods Phys. Res., Sect A 320, 331 (1992).

[6] See /home/lv59/conway/bayes.f.