MEFE 2

Estadística: un campo activo de investigación para físicos

Lo que vamos a ver, en buena parte, surge de papers de estadística publicados por físicos en los últimos 20-30 años.

Nuclear Instruments and Methods in Physics Research 221 (1984) 437–442
North-Holland, Amsterdam

CLARIFICATION OF THE USE OF CHI-SQUARE AND LIKELIHOOD FUNCTIONS IN FITS TO HISTOGRAMS

Steve BAKER and Robert D. COUSINS

Department of Physics, University of California, Los Angeles, California 90024, USA

Received 18 July 1983

We consider the problem of fitting curves to histograms in which the data obey multinomial or Poisson statistics. Techniques commonly used by physicists are examined in light of standard results found in the statistics literature. We review the relationship between multinomial and Poisson distributions, and clarify a sufficient condition for equality of the area under the fitted curve and the number of events on the histogram. Following the statisticians, we use the likelihood ratio test to construct a general χ^2 statistic, χ^2_λ, which yields parameter and error estimates identical to those of the method of maximum likelihood. The χ^2_λ statistic is further useful for testing goodness-of-fit since the value of its minimum asymptotically obeys a classical chi-square distribution. One should be aware, however, of the potential for statistical bias, especially when the number of events is small.

1. Introduction

Standard results from the theory of statistics are often overlooked by scientists searching for sophisticated methods to analyze their data. Though good high-level statistics books have been written expressly for experimentalists [1], we believe that some misunderstanding remains in the professional physics literature.

the language of the statisticians, these are known as (a) point estimation, (b) confidence interval estimation, and (c) goodness-of-fit testing. Because chi-square statistics are sometimes used for all three tasks, the distinction among them can become blurred. However, it is important to maintain this distinction. In general, one need not use the same statistic for all three purposes. Indeed, some of the most powerful tests of goodness-of-

Statistics: an active field for physicists

PHYSICAL REVIEW D VOLUME 57, NUMBER 7 1 APRIL 1998

Unified approach to the classical statistical analysis of small signals

Gary J. Feldman*
Department of Physics, Harvard University, Cambridge, Massachusetts 02138

Robert D. Cousins†
Department of Physics and Astronomy, University of California, Los Angeles, California 90095
(Received 21 November 1997; published 6 March 1998)

We give a classical confidence belt construction which unifies the treatment of upper confidence limits for null results and two-sided confidence intervals for non-null results. The unified treatment solves a problem (apparently not previously recognized) that the choice of upper limit or two-sided intervals leads to intervals which are not confidence intervals if the choice is based on the data. We apply the construction to two related problems which have recently been a battleground between classical and Bayesian statistics: Poisson processes with background and Gaussian errors with a bounded physical region. In contrast with the usual classical construction for upper limits, our construction avoids unphysical confidence intervals. In contrast with some popular Bayesian intervals, our intervals eliminate conservatism (frequentist coverage greater than the stated confidence) in the Gaussian case and reduce it to a level dictated by discreteness in the Poisson case. We generalize the method in order to apply it to analysis of experiments searching for neutrino oscillations. We show that this technique both gives correct coverage and is powerful, while other classical techniques that have been used by neutrino oscillation search experiments fail one or both of these criteria.
[S0556-2821(98)00109-X]

PACS number(s): 06.20.Dk, 14.60.Pq

I. INTRODUCTION

Classical confidence intervals are the traditional way in which high energy physicists report errors on results of experiments. Approximate methods of confidence interval construction, in particular the likelihood ratio method, are often

decide whether to consult confidence interval tables for upper limits or for central confidence intervals. In contrast, our unified set of confidence intervals satisfies (by construction) the classical criterion of frequentist coverage of the unknown true value. Thus the problem of wrong confidence intervals is also solved.

JOURNAL OF MATHEMATICAL PHYSICS VOLUME 41, NUMBER 8 AUGUST 2000

The statistical analysis of Gaussian and Poisson signals near physical boundaries

Mark Mandelkern and Jonas Schultz
Department of Physics and Astronomy, University of California, Irvine, California 92697

(Received 22 February 2000; accepted for publication 17 April 2000)

We propose a construction of frequentist confidence intervals that is effective near unphysical regions and unifies the treatment of two-sided and upper limit intervals. It is rigorous, has coverage, is computationally simple and avoids the pathologies that affect the likelihood ratio and related constructions. Away from nonphysical regions, the results are exactly the usual central two-sided intervals. The construction is based on including the physical constraint in the derivation of the estimator, leading to an estimator with values that are confined to the physical domain.

I. INTRODUCTION

Obtaining confidence intervals near physical boundaries is a long-standing problem. Experiments designed to detect a nonzero neutrino mass by observing neutrino oscillation or to detect a small resonance signal in the presence of background are examples in which a negative result may be obtained for a quantity that is intrinsically positive. The difficulty arises when the estimate for the Gaussian or Poisson mean, as obtained from the data, is near or beyond the physical boundary, in which case the standard (classical) result of Neyman's construction is an unphysical or null interval as illustrated in Figs. 1 and 2.

For the Gaussian case, Fig. 1, one obtains central confidence intervals for the mean μ constrained to be non-negative, using the sample mean $\bar{x}$ as the estimator for μ. $\bar{x}$ sufficiently negative leads to the null interval. Despite the fact that the construction has coverage α, which means that, for any given true mean, the confidence interval includes that value with probability α, the null interval *cannot* contain any true mean. It is necessarily one of the measured intervals that, with probability $1-\alpha$, fail to contain the true mean. Even the non-null intervals obtained by

PHYSICAL REVIEW D **67**, 012002 (2003)

Including systematic uncertainties in confidence interval construction for Poisson statistics

J. Conrad, O. Botner, A. Hallgren, and C. Pérez de los Heros
Division of High Energy Physics, Uppsala University, S-75121 Uppsala, Sweden

(Received 30 January 2002; published 10 January 2003)

One way to incorporate systematic uncertainties into the calculation of confidence intervals is by integrating over probability density functions parametrizing the uncertainties. In this paper we present a development of this method which takes into account uncertainties in the prediction of background processes and uncertainties in the signal detection efficiency and background efficiency, and allows for a correlation between the signal and background detection efficiencies. We implement this method with the likelihood ratio (usually denoted as the Feldman-Cousins) approach with and without conditioning. We present studies of coverage for the likelihood ratio and Neyman ordering schemes. In particular, we present two different types of coverage tests for the case where systematic uncertainties are included. To illustrate the method we show the relative effect of including systematic uncertainties in the case of the dark matter search as performed by modern neutrino telescopes.

DOI: 10.1103/PhysRevD.67.012002 PACS number(s): 06.20.Dk, 95.55.Vj

I. INTRODUCTION

A limit on, or a measurement of, a physical quantity at a given confidence level is usually set by comparing a number of detected events, n_o, with the number of expected events from the known background sources contributing to the physical process in question, n_b. How ''compatible'' these numbers are determines how much room there is for new processes, i.e., for a signal. How well the number of observed events and expected background compare strongly depends on the systematic uncertainties present in the measurement. Systematic uncertainties must, therefore, be taken into account in the limit or confidence belt calculation that is finally published.

Traditionally, confidence limits are set using a Neyman construction [1]. This is a purely frequentist method. Feldman and Cousins [2] have proposed an improved method to

[7], and Durham [8] for a review of the status of the field.

In this paper we extend the method of confidence belt construction proposed in [4] to include systematic uncertainties in both the signal and background efficiencies as well as theoretical uncertainties in the background prediction. The proposed method also allows us to use newer ordering schemes. A recent attempt to include systematic uncertainty in the background prediction in a similar manner has been presented in [9]. The paper is organized as follows. In Sec. II we give a short review of the confidence belt construction schemes that we will use. In Sec. III we describe how to include the systematic uncertainties; in Sec. IV we discuss how the confidence belt construction is performed and present some selected results. We compare the results of this method with other methods to include systematics in Sec. V. We introduce the tests of coverage performed in Sec. VI and present an example based on data from the Antarctic Muon

Available online at www.sciencedirect.com

SCIENCE DIRECT®

ELSEVIER

Nuclear Instruments and Methods in Physics Research A 551 (2005) 493–503

NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A

www.elsevier.com/locate/nima

Limits and confidence intervals in the presence of nuisance parameters

Wolfgang A. Rolke[a,*], Angel M. López[b], Jan Conrad[c]

[a]*Department of Mathematics, University of Puerto Rico - Mayagüez, P.O. Box 5959, Mayagüez, PR 00681, USA*

[b]*Department of Physics, University of Puerto Rico - Mayagüez, Mayagüez, PR 00681, USA*

[c]*PH-Department, CERN, CH-1211, Geneva 23, Switzerland*

Received 7 February 2005; received in revised form 13 May 2005; accepted 18 May 2005
Available online 11 July 2005

Abstract

We study the frequentist properties of confidence intervals computed by the method known to statisticians as the Profile Likelihood. It is seen that the coverage of these intervals is surprisingly good over a wide range of possible parameter values for important classes of problems, in particular whenever there are additional nuisance parameters with statistical or systematic errors. Programs are available for calculating these intervals.

PACS: 06.20.DK; 02.50.−r; 05.10.Ln

PHYSICAL REVIEW D **74,** 092003 (2006)

Signal discovery in sparse spectra: A Bayesian analysis

A. Caldwell and K. Kröninger

Max-Planck-Institut für Physik, München, Föhringer Ring 6, 80805 München, Germany

(Received 19 September 2006; published 20 November 2006)

A Bayesian analysis of the probability of a signal in the presence of background is developed, and criteria are proposed for claiming evidence for, or the discovery of a signal. The method is general and, in particular, applicable to sparsely populated spectra. Monte Carlo techniques to evaluate the sensitivity of an experiment are described. As an example, the method is used to calculate the sensitivity of the GERDA experiment to neutrinoless double beta decay.

DOI: 10.1103/PhysRevD.74.092003 PACS numbers: 02.50.−r, 02.70.Rr, 02.70.Uu, 14.60.St

I. INTRODUCTION

In the analysis of sparsely populated spectra common approximations, valid only for large numbers, fail for the small number of events encountered. A Bayesian analysis of the probability of a signal in the presence of background is developed, and criteria are proposed for claiming evidence for, or the discovery of a signal. It is independent of the physics case and can be applied to a variety of situations.

Model comparisons from a Bayesian perspective have been discussed extensively in the literature [1]. These analyses typically calculate the ‘‘odds‘‘ for one model to be correct relative to the other(s) [2]. In this paper, a somewhat different approach was taken in that a procedure for claiming a discovery is proposed-i.e., for claiming that known processes alone are not enough to describe the measured data.

To make predictions about possible outcomes of an experiment, distributions of quantities under study are calculated. As an approximation, ensembles, sets of Monte Carlo data which mimic the expected spectrum,

signal process to the observed spectrum? What is the probability that the spectrum is due to background only? Given a model for the signal and background, what is the (most probable) parameter value describing the number of signal events in the spectrum? In case no signal is observed, what is the limit that can be set on the signal contribution? The analysis method introduced in this paper is based on Bayes' Theorem and developed to answer these questions and is, in particular, suitable for spectra with a small number of events.

The assumptions for the analysis are

(i) The spectrum is confined to a certain region of interest.

(ii) The spectral shape of a possible signal is known.

(iii) The spectral shape of the background is known [4].

(iv) The spectrum is divided into bins and the event numbers in the bins follow Poisson distributions.

The analysis consists of two steps. First, the probability that the observed spectrum is due to background only is calculated. If this probability is less than an *a priori* defined value, the discovery (or evidence) criterion, the signal

Available online at www.sciencedirect.com

ScienceDirect

Nuclear Instruments and Methods in Physics Research A 570 (2007) 159–164

NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A

www.elsevier.com/locate/nima

Statistical errors in Monte Carlo estimates of systematic errors☆

Byron P. Roe*

Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA

Received 2 August 2006; received in revised form 18 September 2006; accepted 11 October 2006
Available online 3 November 2006

Abstract

For estimating the effects of a number of systematic errors on a data sample, one can generate Monte Carlo (MC) runs with systematic parameters varied and examine the change in the desired observed result. Two methods are often used. In the unisim method,[1] the systematic parameters are varied one at a time by one standard deviation, each parameter corresponding to a MC run. In the multisim method (see footnote 1), each MC run has all of the parameters varied; the amount of variation is chosen from the expected distribution of each systematic parameter, usually assumed to be a normal distribution. The variance of the overall systematic error determination is derived for each of the two methods and comparisons are made between them. If one focuses not on the error in the prediction of an individual systematic error, but on the overall error due to all systematic errors in the error matrix element in data bin m, the number of events needed is strongly reduced because of the averaging effect over all of the errors. For simple models presented here the multisim model was far better if the statistical error in the MC samples was larger than an individual systematic error, while for the reverse case, the unisim model was better. Exact formulas and formulas for the simple toy models are presented so that realistic calculations can be made. The calculations in the present note are valid if the errors are in a linear region. If that region extends sufficiently far, one can have the unisims or multisims correspond to k standard deviations instead of one. This reduces the number of events required by a factor of k^2.

PACS: 2.50.Ng; 2.70.−c; 2.70.Uu; 7.05.Kf

PUBLISHED BY INSTITUTE OF PHYSICS PUBLISHING FOR SISSA

RECEIVED: *April 1, 2006*
REVISED: *December 23, 2006*
ACCEPTED: *December 28, 2006*
PUBLISHED: *January 22, 2007*

Computation of confidence levels for search experiments with fractional event counting and the treatment of systematic errors

Peter Bock

Physikalisches Institut der Universität Heidelberg,
Philosophenweg 12, D69120 Heidelberg, Germany
E-mail: peter.bock@physi.uni-heidelberg.de

ABSTRACT: A method is described which computes, from an observed sample of events, upper limits for the production rate of new particles, or, for the case of an observed excess of events over background, the probability for an upward fluctuation of the background. It is based on weighted event counting depending on a discriminating variable. Candidates may be produced in different reaction channels with different detection efficiencies and different background. Systematic errors with arbitrary correlations are taken into account in the confidence level calculations. In addition, they are are incorporated in the weight definition.

JHEP01(2007)0

Nuclear Instruments and Methods in Physics Research A 595 (2008) 480–501

ELSEVIER

Contents lists available at ScienceDirect

Nuclear Instruments and Methods in Physics Research A

journal homepage: www.elsevier.com/locate/nima

Evaluation of three methods for calculating statistical significance when incorporating a systematic uncertainty into a test of the background-only hypothesis for a Poisson process

Robert D. Cousins [a,*], James T. Linnemann [b], Jordan Tucker [a]

[a] Department of Physics and Astronomy, University of California, Los Angeles, CA 90095, USA
[b] Department of Physics and Astronomy, Michigan State University, East Lansing, MI 48840, USA

ARTICLE INFO

Article history:
Received 25 February 2007
Received in revised form
7 June 2008
Accepted 10 July 2008
Available online 22 July 2008

Keywords:
Hypothesis test
Confidence interval
Systematic uncertainties

ABSTRACT

Hypothesis tests for the presence of new sources of Poisson counts amidst background processes are frequently performed in high energy physics (HEP), gamma ray astronomy (GRA), and other branches of science. While there are conceptual issues already when the mean rate of background is precisely known, the issues are even more difficult when the mean background rate has non-negligible uncertainty. After describing a variety of methods to be found in the HEP and GRA literature, we consider in detail three classes of algorithms and evaluate them over a wide range of parameter space, by the criterion of how close the ensemble-average Type I error rate (rejection of the background-only hypothesis when it is true) compares with the nominal significance level given by the algorithm. We recommend wider use of an algorithm firmly grounded in frequentist tests of the ratio of Poisson means, although for very low counts the overcoverage can be severe due to the effect of discreteness. We extend the studies of Cranmer, who found that a popular Bayesian-frequentist hybrid can undercover severely when taken to high Z-values. We also examine the profile likelihood method, which has long been used in GRA and HEP; it provides an excellent approximation in much of the parameter space, as previously studied by Rolke and collaborators.

Eur. Phys. J. C (2011) 71: 1554
DOI 10.1140/epjc/s10052-011-1554-0

THE EUROPEAN PHYSICAL JOURNAL C

Special Article - Tools for Experiment and Theory

Asymptotic formulae for likelihood-based tests of new physics

Glen Cowan[1], **Kyle Cranmer**[2], **Eilam Gross**[3], **Ofer Vitells**[3,a]

[1]Physics Department, Royal Holloway, University of London, Egham TW20 0EX, UK
[2]Physics Department, New York University, New York, NY 10003, USA
[3]Weizmann Institute of Science, Rehovot 76100, Israel

Received: 15 October 2010 / Revised: 6 January 2011 / Published online: 9 February 2011

Abstract We describe likelihood-based statistical tests for use in high energy physics for the discovery of new phenomena and for construction of confidence intervals on model parameters. We focus on the properties of the test procedures that allow one to account for systematic uncertainties. Explicit formulae for the asymptotic distributions of test statistics are derived using results of Wilks and Wald. We motivate and justify the use of a representative data set, called the "Asimov data set", which provides a simple method to obtain the median experimental sensitivity of a search or measurement as well as fluctuations about this expectation.

1 Introduction

In particle physics experiments one often searches for processes that have been predicted but not yet seen, such as data sets by a single representative one, referred to here as the "Asimov" data set.[1] In the past, this method has been used and justified intuitively (e.g., [4, 5]). Here we provide a formal mathematical justification for the method, explore its limitations, and point out several additional aspects of its use.

The present paper extends what was shown in [5] by giving more accurate formulas for exclusion significance and also by providing a quantitative measure of the statistical fluctuations in discovery significance and exclusion limits. For completeness some of the background material from [5] is summarized here.

In Sect. 2 the formalism of a search as a statistical test is outlined and the concepts of statistical significance and sensitivity are given precise definitions. Several test statistics based on the profile likelihood ratio are defined.

In Sect. 3, we use the approximations due to Wilks and

JINST

PUBLISHED BY IOP PUBLISHING FOR SISSA MEDIALAB

RECEIVED: *December 25, 2014*
REVISED: *February 23, 2015*
ACCEPTED: *March 5, 2015*
PUBLISHED: *April 28, 2015*

Handling uncertainties in background shapes: the discrete profiling method

P.D. Dauncey,[a,1] **M. Kenzie,**[b] **N. Wardle**[b] **and G.J. Davies**[a]

[a] *Department of Physics, Imperial College London, Prince Consort Road, London, SW7 2AZ, U.K.*
[b] *CERN, CH-1211 Geneva 23, Switzerland*

E-mail: P.Dauncey@imperial.ac.uk

ABSTRACT: A common problem in data analysis is that the functional form, as well as the parame-

2015 JINST 10

Astroparticle Physics 62 (2015) 165–177

Contents lists available at ScienceDirect

Astroparticle Physics

journal homepage: www.elsevier.com/locate/astropart

Review

Statistical issues in astrophysical searches for particle dark matter

Jan Conrad

Oskar Klein Centre for Cosmoparticle Physics, Physics Department, Stockholm University, Albanova University Centre, SE-10691 Stockholm and Imperial College London, London SW7 2AZ, UK

A R T I C L E I N F O

Article history:
Received 15 June 2014
Received in revised form 15 August 2014
Accepted 16 September 2014
Available online 27 September 2014

Keywords:
Dark matter
Statistical methods
Indirect detection
Direct detection

A B S T R A C T

In this review statistical issues appearing in astrophysical searches for particle dark matter, i.e. indirect detection (dark matter annihilating into standard model particles) or direct detection (dark matter particles scattering in deep underground detectors) are discussed. One particular aspect of these searches is the presence of very large uncertainties in nuisance parameters (astrophysical factors) that are degenerate with parameters of interest (mass and annihilation/decay cross sections for the particles). The likelihood approach has become the most powerful tool, offering at least one well motivated method for incorporation of nuisance parameters and increasing the sensitivity of experiments by allowing a combination of targets superior to the more traditional data stacking. Other statistical challenges appearing in astrophysical searches are to large extent similar to any new physics search, for example at colliders, a prime example being the calculation of trial factors. Frequentist methods prevail for hypothesis testing and interval estimation, Bayesian methods are used for assessment of nuisance parameters and parameter estimation in complex parameter spaces. The basic statistical concepts will be exposed, illustrated with concrete examples from experimental searches and caveats will be pointed out.

Contents

Estadística, el problema inverso de la probabilidad

Por ejemplo, la distribución exponencial.

$$\boxed{f(x\,;\tau) = \frac{1}{\tau}\,e^{-x/\tau}}$$

Probabilidad : conociendo τ, decir algo de x.

Estadística : conociendo x, decir algo de τ.

Estadística, el problema inverso de la probabilidad

Probabilidad (fácil):

- sabiendo que la vida media de un núcleo es $\tau = 2\,$s
- cual es la fracción de decaimientos entre 1 y $3\,$s?
- Única solución: $P(1 < t < 3 \,|\, \tau = 2) = \int_1^3 f(x \,|\, \tau)\,\mathrm{d}x = e^{-\frac{1}{2}} - e^{-\frac{3}{2}} = 0.3834$

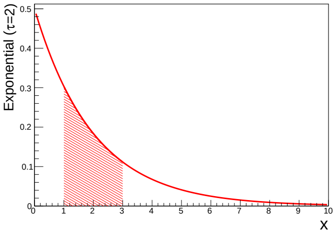

Estadística, el problema inverso de la probabilidad

Estadística (difícil):

- Se hace una medición al azar de la distribución exponencial y da : $x = 2.7$
- Qué podemos decir del parámetro τ ? $\tau = 2.7^{+1.2}_{-0.8}$ with 68% CL.
- La solución es “correcta”, pero no única: muchas otras soluciones correctas.
- A menudo no hay una solución“buena”: muchas soluciones aproximadas.

Herramientas estadísticas en funcionamiento

Make 40 “measurements”: from the exponential $\tau = 2$ distribution, $f(x) = \frac{1}{2}\, e^{-x/2}$, generate 40 random numbers.

```
1.96 0.45 0.10 3.39 0.62 0.50 1.03 1.03 3.72 1.62 3.72 4.61 5.69 1.56 0.29 2.19 1.02 3.89
0.69 0.86 7.77 5.75 0.06 0.65 0.37 0.13 3.10 14.32 3.41 0.02 3.39 0.96 2.38 0.98 1.29 9.90
1.68 7.54 2.47 1.29
```

Herramientas estadísticas en funcionamiento

Make 40 “measurements”: from the exponential $\tau = 2$ distribution, $f(x) = \frac{1}{2}\, e^{-x/2}$, generate 40 random numbers.

```
1.96 0.45 0.10 3.39 0.62 0.50 1.03 1.03 3.72 1.62 3.72 4.61 5.69 1.56 0.29 2.19 1.02 3.89
0.69 0.86 7.77 5.75 0.06 0.65 0.37 0.13 3.10 14.32 3.41 0.02 3.39 0.96 2.38 0.98 1.29 9.90
1.68 7.54 2.47 1.29
```

What to do with these? ☹

Herramientas estadísticas en funcionamiento

Make 40 “measurements”: from the exponential $\tau = 2$ distribution, $f(x) = \frac{1}{2}\, e^{-x/2}$, generate 40 random numbers.

```
1.96 0.45 0.10 3.39 0.62 0.50 1.03 1.03 3.72 1.62 3.72 4.61 5.69 1.56 0.29 2.19 1.02 3.89
0.69 0.86 7.77 5.75 0.06 0.65 0.37 0.13 3.10 14.32 3.41 0.02 3.39 0.96 2.38 0.98 1.29 9.90
1.68 7.54 2.47 1.29
```

What to do with these? ☹

Fill a histogram and fit with exponential

```
r = ROOT.TRandom(0)
h1 = ROOT.TH1D ("h1","Expo",20,0,20)
for x in range(0,40):
   x = r.Exp(2)
   h1.Fill(x)

h1.Fit("expo")
```

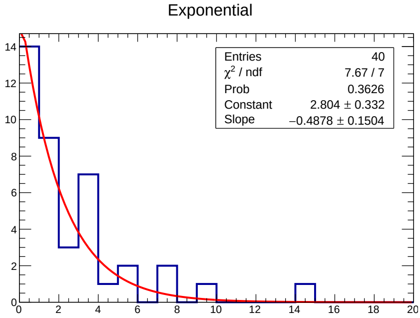

Herramientas estadísticas en funcionamiento

Ya este simple resultado involucra las tres principales áreas de la estadística.

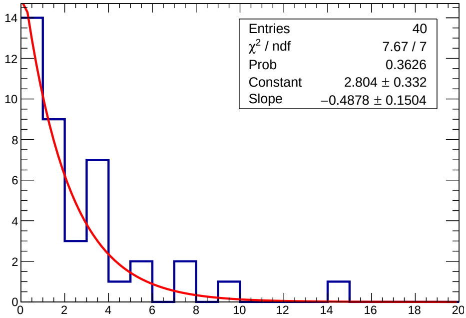

Herramientas estadísticas en funcionamiento

Ya este simple resultado involucra las tres principales áreas de la estadística.

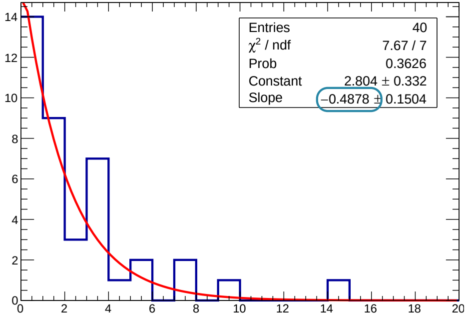

Cual es el valor del parámetro $\hat{\tau}$? $\rightarrow$ Teoría de Estimadores

Herramientas estadísticas en funcionamiento

Ya este simple resultado involucra las tres principales áreas de la estadística.

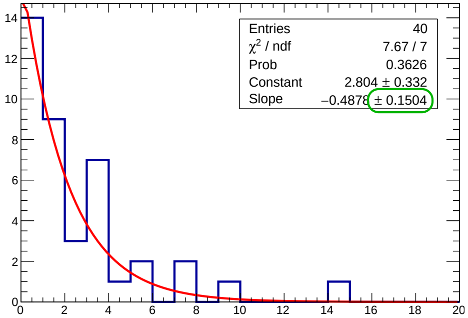

Cual es el valor del parámetro $\hat{\tau}$? $\rightarrow$ Teoría de Estimadores

Cuanto es el error (incerteza) en $\hat{\sigma}_\tau$? $\rightarrow$ Intervalos de Confianza

Herramientas estadísticas en funcionamiento

Ya este simple resultado involucra las tres principales áreas de la estadística.

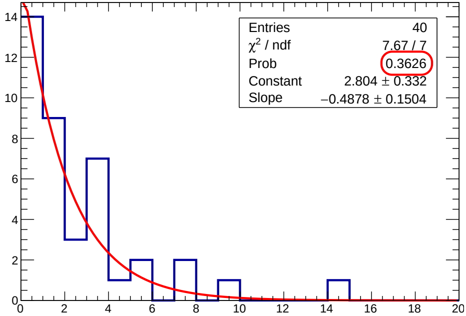

Cual es el valor del parámetro $\hat{\tau}$? $\rightarrow$ Teoría de Estimadores

Cuanto es el error (incerteza) en $\hat{\sigma}_\tau$? $\rightarrow$ Intervalos de Confianza

El fit es bueno? $\rightarrow$ Tests de Hipótesis

Herramientas estadísticas en funcionamiento

Man page de la clase **TH1**:

```
TFitResultPtr TH1::Fit ( TF1 *       f1,
                         Option_t *  option = "",
                         Option_t *  goption = "",
                         Double_t    xxmin = 0,
                         Double_t    xxmax = 0
                       )
```

Fit histogram with function f1.

Parameters

[in] **option** fit options is given in parameter option.

- "W" Set all weights to 1 for non empty bins; ignore error bars
- "WW" Set all weights to 1 including empty bins; ignore error bars
- "I" Use integral of function in bin, normalized by the bin volume, instead of value at bin center
- "L" Use Loglikelihood method (default is chisquare method)
- "WL" Use Loglikelihood method and bin contents are not integer, i.e. histogram is weighted (must have **Sumw2()** set)
- "P" Use Pearson chi2 (using expected errors instead of observed errors)
- "U" Use a User specified fitting algorithm (via SetFCN)
- "Q" Quiet mode (minimum printing)
- "V" Verbose mode (default is between Q and V)
- "E" Perform better Errors estimation using Minos technique
- "B" User defined parameter settings are used for predefined functions like "gaus", "expo", "poln", "landau". Use this opti
- "M" More. Improve fit results. It uses the IMPROVE command of **TMinuit** (see **TMinuit::mnimpr**). This algorithm attem
- "R" Use the Range specified in the function range
- "N" Do not store the graphics function, do not draw
- "0" Do not plot the result of the fit. By default the fitted function is drawn unless the option"N" above is specified.
- "+" Add this new fitted function to the list of fitted functions (by default, any previous function is deleted)
- "C" In case of linear fitting, don't calculate the chisquare (saves time)
- "F" If fitting a polN, switch to minuit fitter
- "S" The result of the fit is returned in the **TFitResultPtr** (see below Access to the Fit Result)

Herramientas estadísticas en funcionamiento

Man page de la clase **TH1**: dos opciones de fiteo adicionales

```
TFitResultPtr TH1::Fit ( TF1 *       f1,
                         Option_t *  option = "",
                         Option_t *  goption = "",
                         Double_t    xxmin = 0,
                         Double_t    xxmax = 0
                       )
```

Fit histogram with function f1.

Parameters

[in] **option** fit options is given in parameter option.

- "W" Set all weights to 1 for non empty bins; ignore error bars
- "WW" Set all weights to 1 including empty bins; ignore error bars
- "I" Use integral of function in bin, normalized by the bin volume, instead of value at bin center
- "L" Use Loglikelihood method (default is chisquare method)
- "WL" Use Loglikelihood method and bin contents are not integer, i.e. histogram is weighted (must have Sumw2() set)
- "P" Use Pearson chi2 (using expected errors instead of observed errors)
- "U" Use a User specified fitting algorithm (via SetFCN)
- "Q" Quiet mode (minimum printing)
- "V" Verbose mode (default is between Q and V)
- "E" Perform better Errors estimation using Minos technique
- "B" User defined parameter settings are used for predefined functions like "gaus", "expo", "poln", "landau". Use this opti
- "M" More. Improve fit results. It uses the IMPROVE command of TMinuit (see TMinuit::mnimpr). This algorithm attem
- "R" Use the Range specified in the function range
- "N" Do not store the graphics function, do not draw
- "0" Do not plot the result of the fit. By default the fitted function is drawn unless the option"N" above is specified.
- "+" Add this new fitted function to the list of fitted functions (by default, any previous function is deleted)
- "C" In case of linear fitting, don't calculate the chisquare (saves time)
- "F" If fitting a polN, switch to minuit fitter
- "S" The result of the fit is returned in the TFitResultPtr (see below Access to the Fit Result)

Herramientas estadísticas en funcionamiento

Man page de la clase **TH1**: dos opciones de fiteo adicionales

```
TFitResultPtr TH1::Fit ( TF1 *       f1,
                         Option_t *  option = "",
                         Option_t *  goption = "",
                         Double_t    xxmin = 0,
                         Double_t    xxmax = 0
                       )
```

Fit histogram with function f1.

Parameters

[in] **option** fit options is given in parameter option.

- "W" Set all weights to 1 for non empty bins; ignore error bars
- "WW" Set all weights to 1 including empty bins; ignore error bars
- "I" Use integral of function in bin, normalized by the bin volume, instead of value at bin center
- "L" Use Loglikelihood method (default is chisquare method)
- "WL" Use Loglikelihood method and bin contents are not integer, i.e. histogram is weighted (must have **Sumw2()** set)
- "P" Use Pearson chi2 (using expected errors instead of observed errors)
- "U" Use a User specified fitting algorithm (via SetFCN)
- "Q" Quiet mode (minimum printing)
- "V" Verbose mode (default is between Q and V)
- "E" Perform better Errors estimation using Minos technique
- "B" User defined parameter settings are used for predefined functions like "gaus", "expo", "poln", "landau". Use this opti
- "M" More. Improve fit results. It uses the IMPROVE command of **TMinuit** (see **TMinuit::mnimpr**). This algorithm attem
- "R" Use the Range specified in the function range
- "N" Do not store the graphics function, do not draw
- "0" Do not plot the result of the fit. By default the fitted function is drawn unless the option"N" above is specified.
- "+" Add this new fitted function to the list of fitted functions (by default, any previous function is deleted)
- "C" In case of linear fitting, don't calculate the chisquare (saves time)
- "F" If fitting a polN, switch to minuit fitter
- "S" The result of the fit is returned in the **TFitResultPtr** (see below Access to the Fit Result)

Herramientas estadísticas en funcionamiento

Generate 40 random numbers from the exponential $f(x) = \frac{1}{2}\,e^{-x/2}$

Fit them using the three alternatives:

```
r = ROOT.TRandom(0)
h1 = ROOT.TH1D ("h1","Expo",20,0,20)
for x in range(0,40):
    x = r.Exp(2)
    h1.Fill(x)

h1.Fit("expo")
```

Herramientas estadísticas en funcionamiento

Generate 40 random numbers from the exponential $f(x) = \frac{1}{2}\, e^{-x/2}$

Fit them using the three alternatives:

```
r = ROOT.TRandom(0)
h1 = ROOT.TH1D ("h1","Expo",20,0,20)
for x in range(0,40):
    x = r.Exp(2)
    h1.Fill(x)

h1.Fit("expo")
h1.Fit("expo","P+")
```

Herramientas estadísticas en funcionamiento

Generate 40 random numbers from the exponential $f(x) = \frac{1}{2}\, e^{-x/2}$

Fit them using the three alternatives:

```
r = ROOT.TRandom(0)
h1 = ROOT.TH1D ("h1","Expo",20,0,20)
for x in range(0,40):
    x = r.Exp(2)
    h1.Fill(x)

h1.Fit("expo")
h1.Fit("expo","P+")
h1.Fit("expo","L+")
```

Herramientas estadísticas en funcionamiento

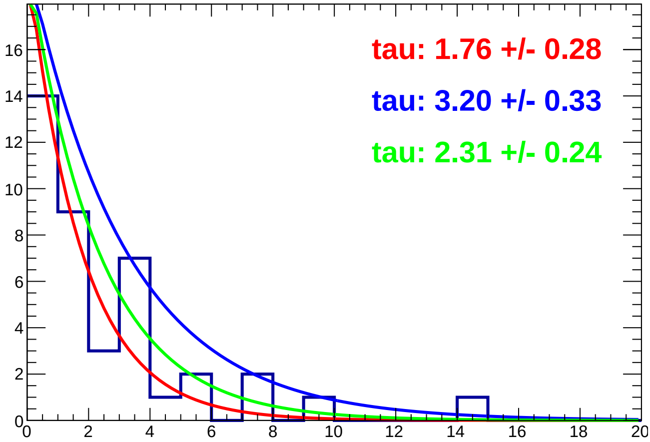

Neyman Pearson LLikelihood

Herramientas estadísticas en funcionamiento

- Así que tres resultados diferentes!
- Cual es el correcto? Respuesta: *NINGUNO*
- Curiosamente, la respuesta podría haber sido: *LOS TRES*
- Pero, que significa “el correcto”?

Herramientas estadísticas en funcionamiento

- Así que tres resultados diferentes!
- Cual es el correcto? Respuesta: *NINGUNO*
- Curiosamente, la respuesta podría haber sido: *LOS TRES*
- Pero, que significa "el correcto"?
- **That is the question.**

Herramientas estadísticas en funcionamiento

Para eso hacemos pseudo-experimentos:

Repetimos muchas veces el experimento de extraer 40 números al azar de la distribución Exp(2), hacemos los fits N P L, y anotamos:

1. los resultados estimados de los parámetros, $\hat{\tau}_N$, $\hat{\tau}_P$, $\hat{\tau}_L$.

Herramientas estadísticas en funcionamiento

Para eso hacemos pseudo-experimentos:

Repetimos muchas veces el experimento de extraer 40 números al azar de la distribución Exp(2), hacemos los fits N P L, y anotamos:

1. los resultados estimados de los parámetros, $\hat{\tau}_N$, $\hat{\tau}_P$, $\hat{\tau}_L$.
2. si el rango $\hat{\tau} \pm \hat{\sigma}_\tau$ incluye o no el verdadero valor $\tau = 2$.

Herramientas estadísticas en funcionamiento

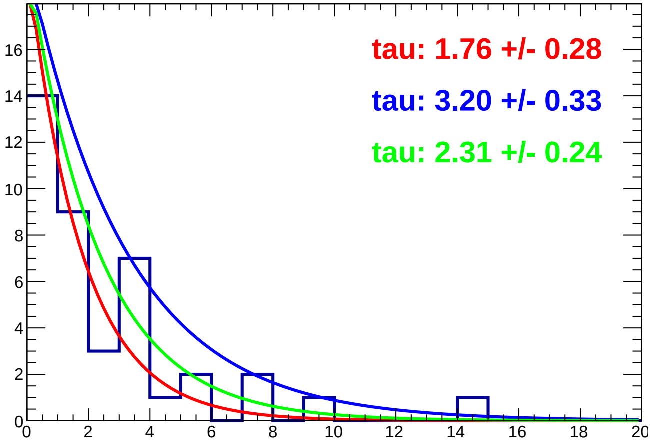

Neyman Pearson LLikelihood

Herramientas estadísticas en funcionamiento

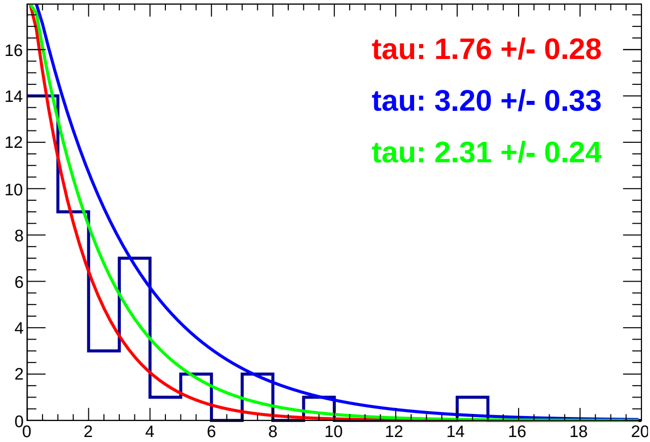

Neyman ☺ Pearson ☹ LLikelihood ☹

Herramientas estadísticas en funcionamiento

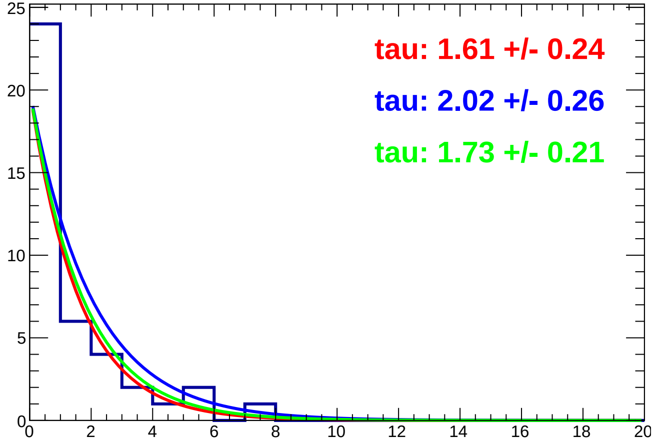

Neyman ☹ Pearson ☺ LLikelihood ☹

Herramientas estadísticas en funcionamiento

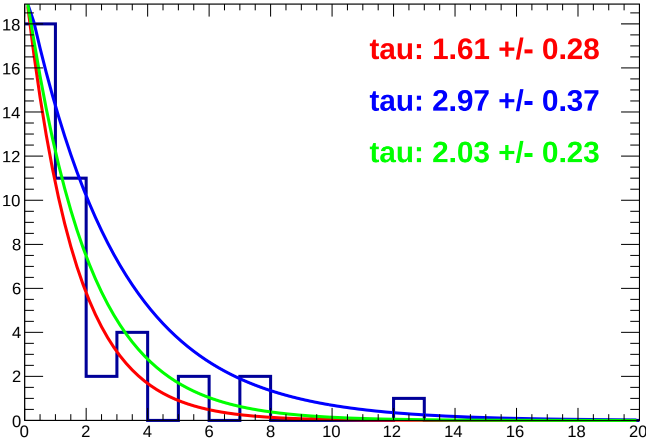

Neyman ☹ Pearson ☹ LLikelihood ☺

Herramientas estadísticas en funcionamiento

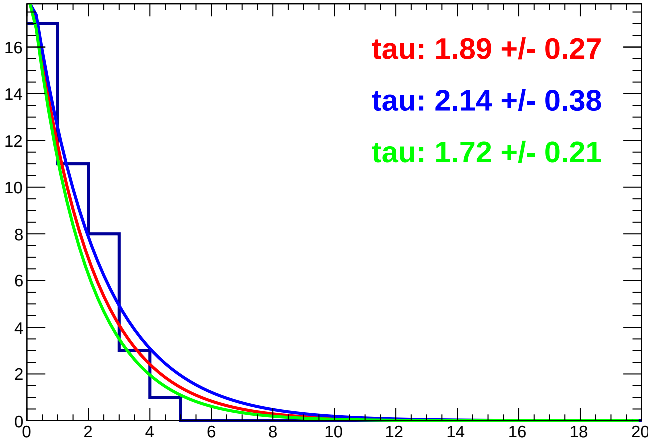

Neyman ☺ Pearson ☺ LLikelihood ☹

Herramientas estadísticas en funcionamiento

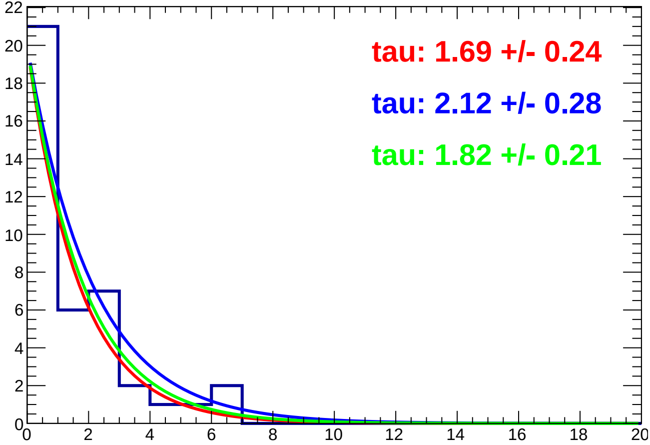

Neyman ☹ Pearson ☺ LLikelihood ☺

Herramientas estadísticas en funcionamiento

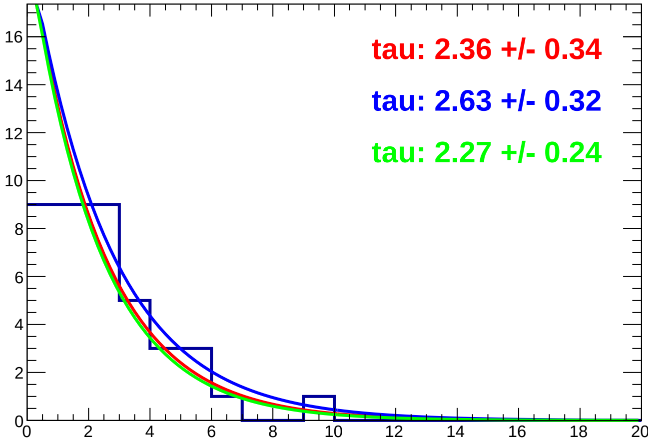

Neyman ☹ Pearson ☹ LLikelihood ☹

Herramientas estadísticas en funcionamiento

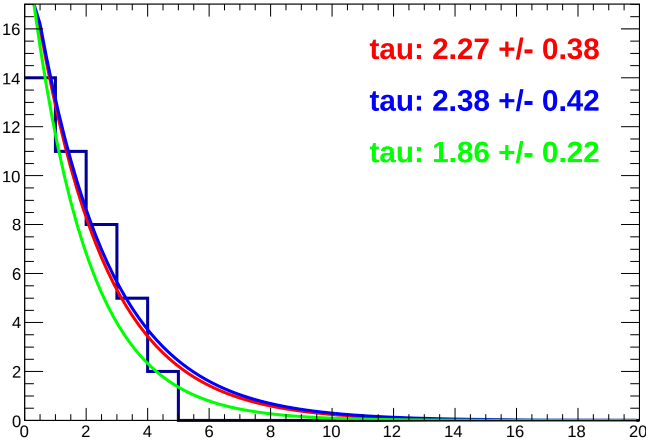

Neyman ☺ Pearson ☺ LLikelihood ☺

Binned vs Unbinned Fits

- The histogram fit uses the number of entries per bin.
- The information of each individual value is lost $\Rightarrow$ BINNED FIT.
- Another possibility is to fit the 40 numbers themselves $\Rightarrow$ UNBINNED FIT.

Binned vs Unbinned Fits

Binned Fit

```
r = ROOT.TRandom(0)
h1 = ROOT.TH1D("h1","Expo",20,0,20)
fE = ROOT.TF1("fE","[a]*exp(-x/[b])",0,20)
for i in range(0,40):
   x = r.Exp(2)
   h1.Fill(x)


fE.SetParameters(2.,2.)
h1.Fit("fE")
```

Unbinned Fit

```
r = ROOT.TRandom(0)
t1 = ROOT.TNtuple("t1","Expo","x")
fE = ROOT.TF1("fE","(1/[b])*exp(-x/[b])",0,20)
for i in range(0,40):
   x = r.Exp(2)
   t1.Fill(x)


fE.SetParameter("b",2.)
t1.UnbinnedFit("fE","x");
```

Intervalos de Confianza: la aproximación parabólica

Desarrollemos la verosimilitud $\log \mathscr{L}(\theta)$ en serie alrededor de $\hat{\theta}$

$$\log \mathscr{L}(\theta) = \log \mathscr{L}(\hat{\theta}) + \left.\frac{\partial \log \mathscr{L}(\theta)}{\partial \theta}\right|_{\hat{\theta}} (\theta - \hat{\theta}) + \frac{1}{2} \left.\frac{\partial^2 \log \mathscr{L}}{\partial \theta^2}\right|_{\hat{\theta}} (\theta - \hat{\theta})^2$$

En el término lineal, la derivada es cero.

En el cuadrático, es el inverso del estimador Cramer-Rao de sigma, $\hat{\sigma}^2 = -\left.\frac{\partial^2 \log \mathscr{L}}{\partial \theta^2}\right|_{\hat{\theta}}^{-1}$

$$\log \mathscr{L}(\theta) = \log \mathscr{L}_{\text{max}} - \frac{1}{2} \left(\frac{\theta - \hat{\theta}}{\hat{\sigma}} \right)^2$$

En la proximidad del estimador ML, $\log \mathscr{L}$ es parabólica.

Intervalos de Confianza: la aproximación parabólica

Para ver en que rango vale la aproximación, es útil la interpretación gráfica de σ.

$$\log \mathscr{L}_{\mathrm{p}}(\theta) = \log \mathscr{L}_{\mathsf{max}} - \frac{1}{2}\left(\frac{\theta - \hat{\theta}}{\sigma}\right)^2$$

$$\log \mathscr{L}_{\mathrm{p}}(\theta = \hat{\theta} \pm n\sigma) = \log \mathscr{L}_{\mathsf{max}} - \frac{n^2}{2}$$

$$\log \mathscr{L}_{\mathrm{p}}(\hat{\theta} \pm n\sigma) = \log \mathscr{L}_{\mathrm{max}} - \frac{n^2}{2}$$

Intervalos de Confianza: la aproximación parabólica

Para ver en que rango vale la aproximación, es útil la interpretación gráfica de σ.

$$\log \mathscr{L}_{\mathrm{p}}(\theta) = \log \mathscr{L}_{\mathsf{max}} - \frac{1}{2}\left(\frac{\theta - \hat{\theta}}{\sigma}\right)^2$$

$$\log \mathscr{L}_{\mathrm{p}}(\theta = \hat{\theta} \pm n\sigma) = \log \mathscr{L}_{\mathsf{max}} - \frac{n^2}{2}$$

$$\log \mathscr{L}_{\mathrm{p}}(\hat{\theta} \pm n\sigma) = \log \mathscr{L}_{\mathrm{max}} - \frac{n^2}{2}$$

$$68.3\%\, CL \;\; (\hat{\theta} \pm 1\sigma) \; : \quad \log \mathscr{L}_{\mathrm{p}}(\theta_{\mathrm{CI}}) = \log \mathscr{L}_{\mathrm{max}} - \frac{1^2}{2} \quad (0.5)$$

$$95.4\%\, CL \;\; (\hat{\theta} \pm 2\sigma) \; : \quad \log \mathscr{L}_{\mathrm{p}}(\theta_{\mathrm{CI}}) = \log \mathscr{L}_{\mathrm{max}} - \frac{2^2}{2} \quad (2.0)$$

$$99.7\%\, CL \;\; (\hat{\theta} \pm 3\sigma) \; : \quad \log \mathscr{L}_{\mathrm{p}}(\theta_{\mathrm{CI}}) = \log \mathscr{L}_{\mathrm{max}} - \frac{3^2}{2} \quad (4.5)$$

Intervalos de Confianza: la aproximación parabólica

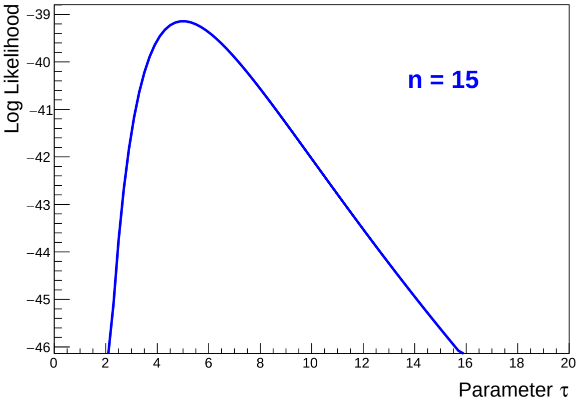

The log-likelihood for $n = 15$ measurements from an exponential distribution

$$\mathscr{L}(\tau \,|\, \boldsymbol{x}) = \frac{1}{\tau^n}\, e^{-\sum x_i/\tau}$$

Intervalos de Confianza: la aproximación parabólica

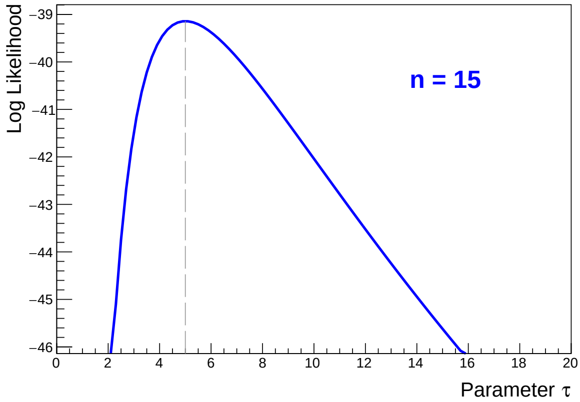

The log-likelihood for $n = 15$ measurements from an exponential distribution

$\mathscr{L}(\tau \mid \boldsymbol{x}) = \dfrac{1}{\tau^n}\, e^{-\sum x_i/\tau}$ The ML estimate in this example is $\hat{\tau} = \dfrac{\sum x_i}{n} = 5.0$

Intervalos de Confianza: la aproximación parabólica

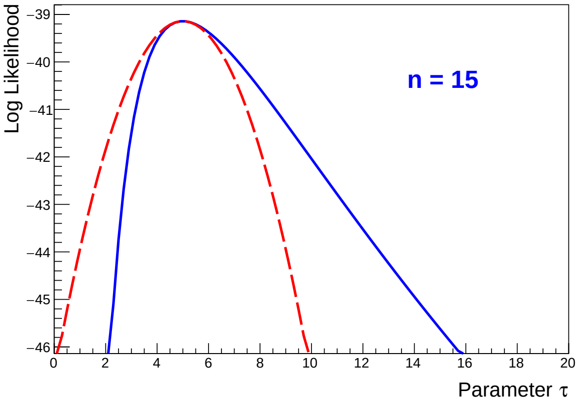

Trace the parabolic approximation $\mathscr{L}_{\mathrm{p}}$ that matches $\mathscr{L}$ at the maximum $(\mathscr{L}_{\max})$

Intervalos de Confianza: la aproximación parabólica

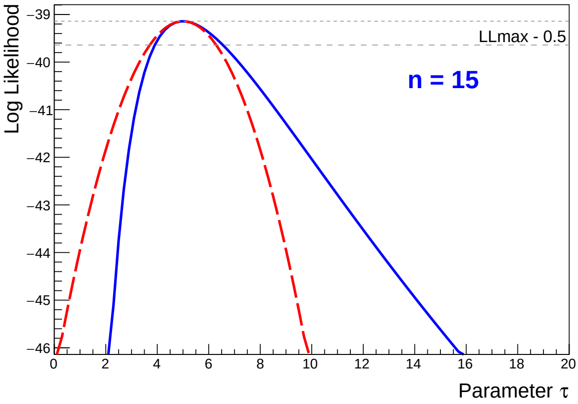

Draw a line at $\mathscr{L}_{\max} - 0.5$

Intervalos de Confianza: la aproximación parabólica

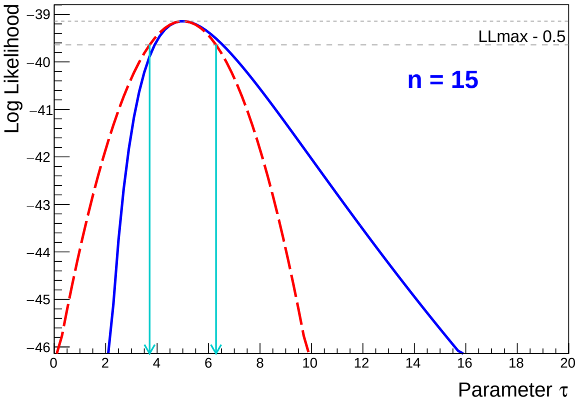

The 68.3% CL interval limits correspond to the points where $\mathscr{L}_{\mathrm{p}} = \mathscr{L}_{\max} - 0.5$

Intervalos de Confianza: la aproximación parabólica

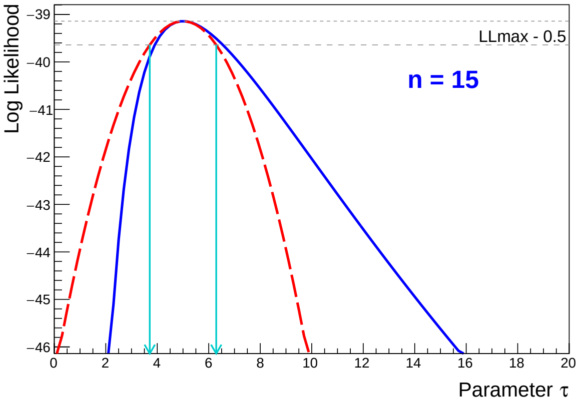

$\log \mathscr{L}_{\mathrm{p}}(\tau_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 0.5 \qquad 68.3\%\, CL \;\; (\hat{\tau} \pm 1\sigma) : \quad 5.0 \pm 1.3$

Intervalos de Confianza: la aproximación parabólica

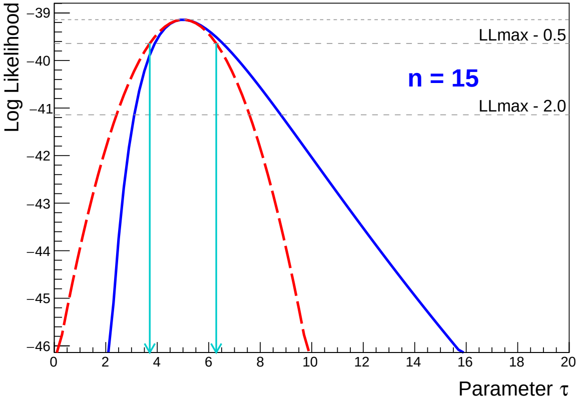

$\log \mathscr{L}_{\mathrm{p}}(\tau_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 0.5 \qquad 68.3\%\, CL \;\; (\hat{\tau} \pm 1\sigma): \quad 5.0 \pm 1.3$

Intervalos de Confianza: la aproximación parabólica

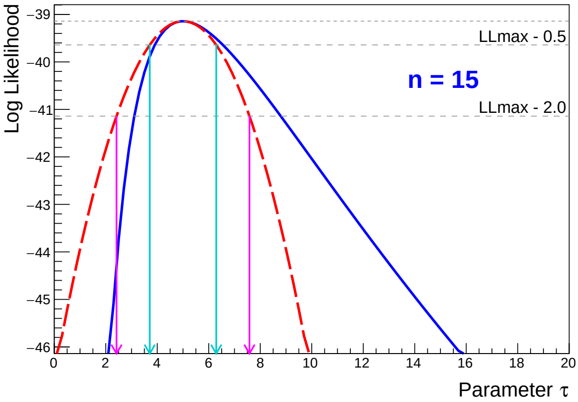

$$\log \mathscr{L}_{\rm p}(\tau_{\rm CI}) = \log \mathscr{L}_{\rm max} - 0.5 \qquad 68.3\%\, CL \quad (\hat{\tau} \pm 1\sigma): \quad 5.0 \pm 1.3$$

$$\log \mathscr{L}_{\rm p}(\tau_{\rm CI}) = \log \mathscr{L}_{\rm max} - 2.0 \qquad 95.4\%\, CL \quad (\hat{\tau} \pm 2\sigma): \quad 5.0 \pm 2.6$$

Intervalos de Confianza: la aproximación parabólica

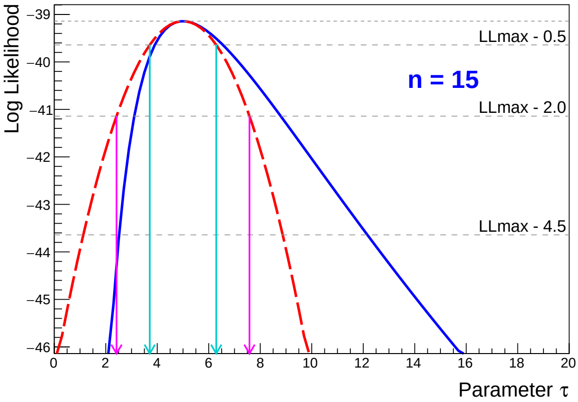

$$\log \mathscr{L}_{\mathrm{p}}(\tau_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 0.5 \qquad 68.3\%\, CL \;\; (\hat{\tau} \pm 1\sigma): \quad 5.0 \pm 1.3$$

$$\log \mathscr{L}_{\mathrm{p}}(\tau_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 2.0 \qquad 95.4\%\, CL \;\; (\hat{\tau} \pm 2\sigma): \quad 5.0 \pm 2.6$$

Intervalos de Confianza: la aproximación parabólica

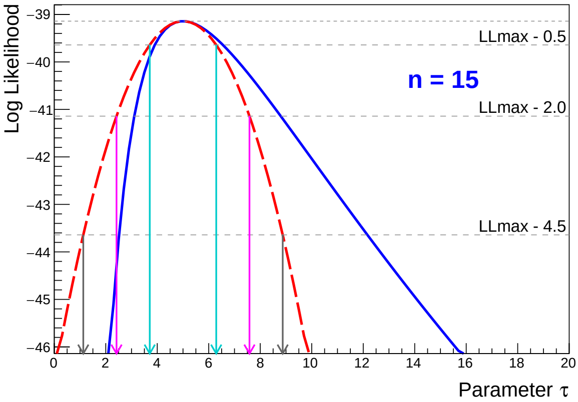

$$\log \mathscr{L}_{\mathrm{p}}(\tau_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 0.5 \qquad 68.3\%\, CL \quad (\hat{\tau} \pm 1\sigma): \quad 5.0 \pm 1.3$$

$$\log \mathscr{L}_{\mathrm{p}}(\tau_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 2.0 \qquad 95.4\%\, CL \quad (\hat{\tau} \pm 2\sigma): \quad 5.0 \pm 2.6$$

$$\log \mathscr{L}_{\mathrm{p}}(\tau_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 4.5 \qquad 99.7\%\, CL \quad (\hat{\tau} \pm 3\sigma): \quad 5.0 \pm 3.9$$

Intervalos de Confianza: la aproximación parabólica

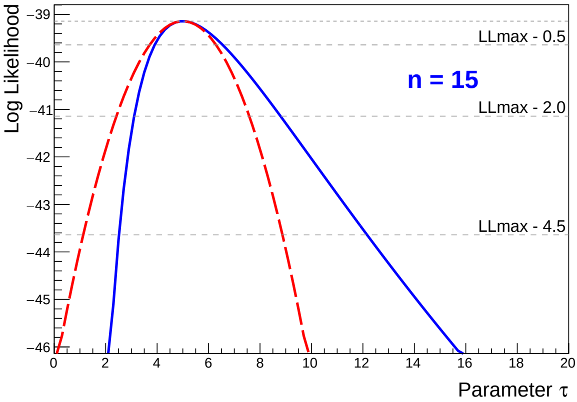

For $n = 15$, the parabolic approximation works fine up to $\pm 1\sigma$

Intervalos de Confianza: la aproximación parabólica

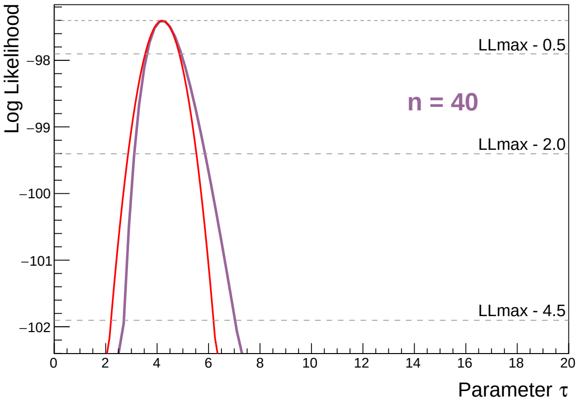

For $n = 40$, the parabolic approximation should work fine for $\pm 1\sigma$, $\pm 2\sigma$, and $\pm 3\sigma$

Intervalos de Confianza: la aproximación parabólica

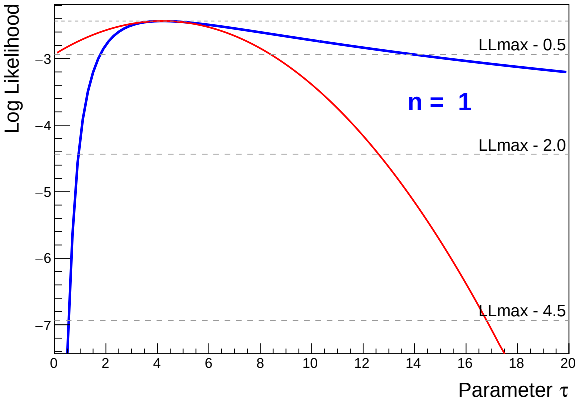

For $n = 1$, the parabolic approximation does not work even for $\pm 1\sigma$

Confidence Intervals: the log-likelihood ratio (LLR) approximation

The recipe for the LLR intervals (also called MINOS confidence intervals) is very simple.

Take the recipe for the parabolic approximation:

$$\begin{aligned}
68.3\%\,CL &: \quad \log \mathscr{L}_{\mathrm{p}}(\theta_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 0.5 \\
95.4\%\,CL &: \quad \log \mathscr{L}_{\mathrm{p}}(\theta_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 2.0 \\
99.7\%\,CL &: \quad \log \mathscr{L}_{\mathrm{p}}(\theta_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 4.5
\end{aligned}$$

Confidence Intervals: the log-likelihood ratio (LLR) approximation

The recipe for the LLR intervals (also called MINOS confidence intervals) is very simple.

Take the recipe for the parabolic approximation:

$$
\begin{aligned}
68.3\%\, CL &: \quad \log \mathscr{L}_{\mathrm{p}}(\theta_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 0.5 \\
95.4\%\, CL &: \quad \log \mathscr{L}_{\mathrm{p}}(\theta_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 2.0 \\
99.7\%\, CL &: \quad \log \mathscr{L}_{\mathrm{p}}(\theta_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 4.5
\end{aligned}
$$

and use the correct likelihood instead of its parabolic approximation:

$$
\begin{aligned}
68.3\%\, CL &: \quad \log \mathscr{L}\ (\theta_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 0.5 \\
95.4\%\, CL &: \quad \log \mathscr{L}\ (\theta_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 2.0 \\
99.7\%\, CL &: \quad \log \mathscr{L}\ (\theta_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 4.5
\end{aligned}
$$

Confidence Intervals: the log-likelihood ratio (LLR) approximation

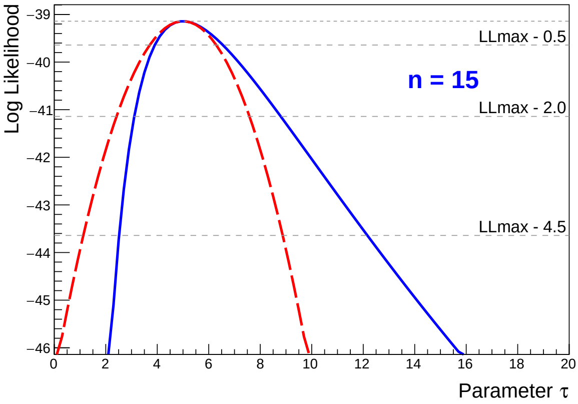

Start with the full likelihood $\mathscr{L}(\tau)$.

Draw lines at $\mathscr{L}_{\max} - \frac{1^2}{2}$, $\mathscr{L}_{\max} - \frac{2^2}{2}$, and $\mathscr{L}_{\max} - \frac{3^2}{2}$.

Confidence Intervals: the log-likelihood ratio (LLR) approximation

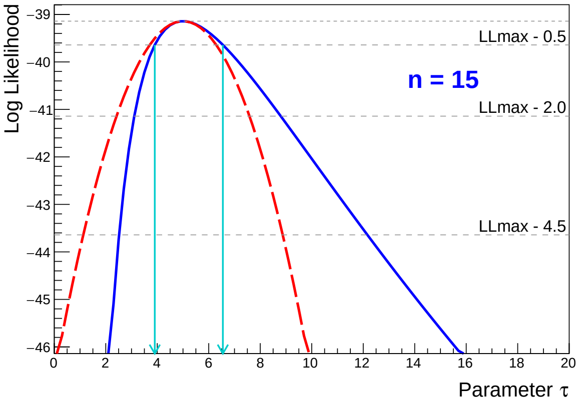

$\log \mathscr{L}(\tau_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 0.5 \qquad 68.3\%\, CL \;:\quad 5.0^{+1.5}_{-1.1} \qquad (5.0 \pm 1.3)$

Confidence Intervals: the log-likelihood ratio (LLR) approximation

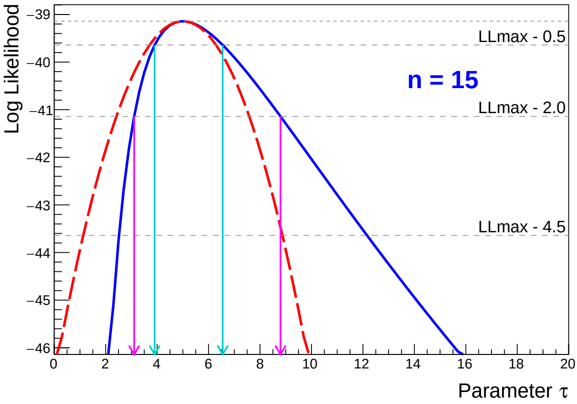

$$\log \mathscr{L}(\tau_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 0.5 \qquad 68.3\%\, CL \;:\quad 5.0^{+1.5}_{-1.1} \qquad (5.0 \pm 1.3)$$

$$\log \mathscr{L}(\tau_{\mathrm{CI}}) = \log \mathscr{L}_{\max} - 2.0 \qquad 95.4\%\, CL \;:\quad 5.0^{+3.8}_{-1.9} \qquad (5.0 \pm 2.6)$$

Confidence Intervals: the log-likelihood ratio (LLR) approximation

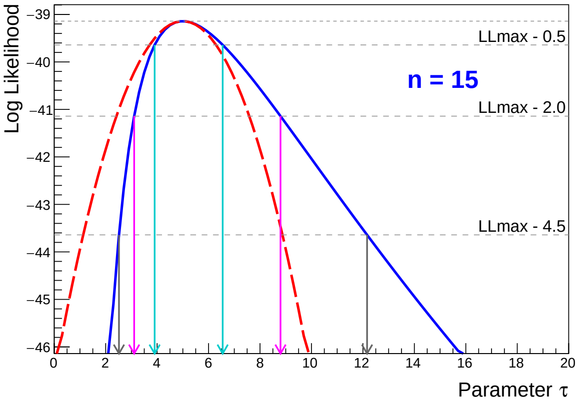

$$\log \mathscr{L}(\tau_{\mathrm{CI}}) = \log \mathscr{L}_{\mathrm{max}} - 0.5 \qquad 68.3\%\,CL \;:\; 5.0^{+1.5}_{-1.1} \qquad (5.0 \pm 1.3)$$

$$\log \mathscr{L}(\tau_{\mathrm{CI}}) = \log \mathscr{L}_{\mathrm{max}} - 2.0 \qquad 95.4\%\,CL \;:\; 5.0^{+3.8}_{-1.9} \qquad (5.0 \pm 2.6)$$

$$\log \mathscr{L}(\tau_{\mathrm{CI}}) = \log \mathscr{L}_{\mathrm{max}} - 4.5 \qquad 99.7\%\,CL \;:\; 5.0^{+7.2}_{-2.5} \qquad (5.0 \pm 3.9)$$

Confidence Intervals: the log-likelihood ratio (LLR) approximation

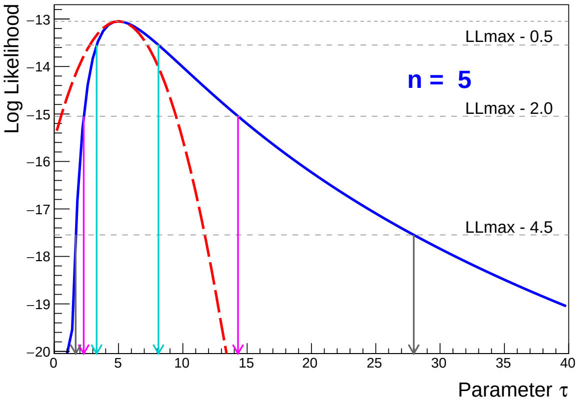

$$\log \mathscr{L}(\tau_{\rm CI}) = \log \mathscr{L}_{\rm max} - 0.5 \qquad 68.3\%\, CL \;:\; 5.0^{+3.1}_{-1.7} \qquad (5.0 \pm 2.2)$$

$$\log \mathscr{L}(\tau_{\rm CI}) = \log \mathscr{L}_{\rm max} - 2.0 \qquad 95.4\%\, CL \;:\; 5.0^{+9.3}_{-2.7} \qquad (5.0 \pm 4.5)$$

$$\log \mathscr{L}(\tau_{\rm CI}) = \log \mathscr{L}_{\rm max} - 4.5 \qquad 99.7\%\, CL \;:\; 5.0^{+23.3}_{-3.5} \qquad (5.0 \pm 6.7)$$

Confidence Intervals: the log-likelihood ratio (LLR) approximation

Assessing the quality of the likelihood intervals:

- Generate a sample of n measurements from the exponential $\tau = 2$ distribution.
- Calculate for this sample the parabolic or LLR interval at a given CL.
- Check if the true value ($\tau = 2$) is included in (covered by) the interval.
- Repeat the experiment many times ($N = 10^6$, for instance).
- Measure coverage: fraction of experiments where CI contains the true value.
- Verify if the coverage of the interval is equal to its Confidence Level.

The Likelihood Function: extracting the Confidence Interval from $\mathscr{L}$

Coverage results for n measurements from an exponential distribution with $\tau = 2$.

The CI are obtained at each CL using the parabolic and the LLR approximations.

	68.3%		95.4%		99.7%	
n	Parab.	LLR	Parab.	LLR	Parab.	LLR
1	60.5	64.4	71.8	80.7	77.1	81.2
2	66.6	66.3	79.8	92.6	85.9	93.6
3	67.7	66.9	83.4	95.1	89.6	97.4
4	68.4	67.2	85.4	94.6	92.1	99.1
5	68.1	68.0	87.0	94.9	93.1	99.5
6	68.2	68.5	87.8	95.2	94.7	99.6
8			90.1	95.3	95.4	99.6
10			90.5	95.0	96.3	99.6
15			92.3	95.5	97.0	99.6
20			92.8	95.2	97.8	99.6
30			93.7	95.4	98.3	99.7
40			93.8	95.2	98.8	99.7