

Statistical analysis of experimental data

Parameter Inference

Aleksander Filip Żarnecki



Lecture 05

November 10, 2022


Parameter Inference

- 1 Maximum Likelihood Method
- 2 Confidence intervals
- 3 Real life example

Applications


Described procedure can be used not only to calculate integrals of one-dimensional functions, it is much more general...

How to calculate volume of a given shape?



Standard procedure:

scan all dimensions using dense point grid and sum cells with centers inside the volume



Monte Carlo integration:

Generate random points in the considered parameter space and count points inside the volume

General case

Examples presented considered the special case: input random variables had uniform distribution and “test function” was binary (returning 0 or 1).

In the general case we want to determine an expectation value of a function $h(\mathbf{x})$ of random variable vector \mathbf{x} described by $f(\mathbf{x})$ pdf:

$$\mu_h \equiv \mathbb{E}_f[h(\mathbf{x})] = \int d\mathbf{x} h(\mathbf{x}) f(\mathbf{x})$$

Monte Carlo determination of μ_h assumes we can generate random variables according to $f(\mathbf{x})$. We can then calculate:

$$\mu_{MC} = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_i h(\mathbf{x}_i)$$

where \mathbf{x}_i , $i = 1, \dots, N$ are random (input) variables generated from $f(\mathbf{x})$

Importance sampling

When $h(\mathbf{x})$ varies strongly in the considered variable range, statistical precision on the mean can be poor. Can it be improved?

Possible solution is to generate \mathbf{x} using probability density more “focused” on the areas where $h(\mathbf{x})$ is large. Optimal choice turns out to be

$$g(\mathbf{x}) \sim h(\mathbf{x}) f(\mathbf{x})$$

but approximate descriptions also work well.

When generating input variables from $g(\mathbf{x})$, the mean value of $h(\mathbf{x})$ can be now calculated as:

$$\mu_{IS} = \frac{1}{N} \sum_i h(\mathbf{x}_i) \cdot \frac{f(\mathbf{x}_i)}{g(\mathbf{x}_i)}$$

where the second term corrects for the modified pdf.

Weighted Monte Carlo

When using weighted Monte Carlo “events”, number of events has to be replaced by sum of weights:

$$N \rightarrow N_w = \sum_i w_i$$

Variance of the sum of weights:

$$\mathbb{V}(N_w) = \sum_i w_i^2$$

Statistical power of the weighted Monte Carlo sample is equivalent to:

$$N_{eq} = \frac{N_w^2}{\mathbb{V}(N_w)} = \frac{(\sum_i w_i)^2}{\sum_i w_i^2}$$

For Poisson distributed random number $\mathbb{V}(N) = N$

Weighted mean

$$\mathbf{I} = (1, \dots, 1)^T$$

What about averaging measurements which are not independent?

In the most general case, variance of the weighted mean is given by

$$\sigma_{\bar{x}}^2 = \mathbf{a}^T \mathbf{C}_x \mathbf{a} - 2 \lambda (\mathbf{a}^T \mathbf{I} - 1)$$

Minimizing mean variance we compare partial derivatives to zero and get

$$\mathbf{C}_x \mathbf{a} = \lambda \cdot \mathbf{I}$$

where λ can be constrained from the boundary condition $\mathbf{a}^T \mathbf{I} = 1$.

This is a linear set of equations, which can be solved:

$$\mathbf{a} = \frac{\mathbf{C}_x^{-1} \mathbf{I}}{\mathbf{I}^T \mathbf{C}_x^{-1} \mathbf{I}}$$

Maximum Likelihood Method

The product:

$$L = \prod_{j=1}^N f(\mathbf{x}^{(j)}; \boldsymbol{\lambda})$$

is called a **likelihood function**.

The most commonly used approach to parameter estimation is the **maximum likelihood approach**:

as the **best estimate of the parameter set $\boldsymbol{\lambda}$** we choose the parameter values for which the **likelihood function has a (global) maximum**.

Frequently used is also log-likelihood function

$$\ell = \ln L = \sum_{j=1}^N \ln f(\mathbf{x}^{(j)}; \boldsymbol{\lambda})$$

we can look for maximum value of ℓ or minimum of $-2\ell = -2\ln L$

Parameter Inference

- 1 Maximum Likelihood Method
- 2 Confidence intervals
- 3 Real life example

Mean estimate

modified from lecture 04

Let us consider N independent measurements of variable X with given (uniform) uncertainty. Assuming measurement fluctuations are described by Gaussian pdf, the likelihood function is:

$$L = \prod_{i=1}^N G(x_i; \mu, \sigma) = \prod_{i=1}^N \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x_i - \mu)^2}{\sigma^2}\right)$$

Log-likelihood:

$$\ell = -\frac{1}{2\sigma^2} \sum (x_i - \mu)^2 + \text{const}$$

$$\frac{\partial \ell}{\partial \mu} = \frac{1}{\sigma^2} \sum (x_i - \mu) = 0$$

$$\Rightarrow \mu = \frac{1}{N} \sum x_i$$

Variance estimate

The same method can be also used to estimate the variance of the Gaussian distribution. We just need to consider partial derivative with respect to σ^2 . Log-likelihood (we can not neglect normalization now):

$$\ell = -\frac{1}{2\sigma^2} \sum (x_i - \mu)^2 - \frac{N}{2} \ln(2\pi\sigma^2)$$

Variance estimate

The same method can be also used to estimate the variance of the Gaussian distribution. We just need to consider partial derivative with respect to σ^2 . Log-likelihood (we can not neglect normalization now):

$$\ell = -\frac{1}{2\sigma^2} \sum (x_i - \mu)^2 - \frac{N}{2} \ln(2\pi\sigma^2)$$

$$\frac{\partial \ell}{\partial \sigma^2} = \frac{1}{2\sigma^4} \sum (x_i - \mu)^2 - \frac{N}{2\sigma^2}$$

Variance estimate

The same method can be also used to estimate the variance of the Gaussian distribution. We just need to consider partial derivative with respect to σ^2 . Log-likelihood (we can not neglect normalization now):

$$\ell = -\frac{1}{2\sigma^2} \sum (x_i - \mu)^2 - \frac{N}{2} \ln(2\pi\sigma^2)$$

$$\frac{\partial \ell}{\partial \sigma^2} = \frac{1}{2\sigma^4} \sum (x_i - \mu)^2 - \frac{N}{2\sigma^2} = 0$$

$$\Rightarrow \sigma^2 = \frac{1}{N} \sum (x_i - \mu)^2$$

Variance estimate

The same method can be also used to estimate the variance of the Gaussian distribution. We just need to consider partial derivative with respect to σ^2 . Log-likelihood (we can not neglect normalization now):

$$\ell = -\frac{1}{2\sigma^2} \sum (x_i - \mu)^2 - \frac{N}{2} \ln(2\pi\sigma^2)$$

$$\frac{\partial \ell}{\partial \sigma^2} = \frac{1}{2\sigma^4} \sum (x_i - \mu)^2 - \frac{N}{2\sigma^2} = 0$$

$$\Rightarrow \sigma^2 = \frac{1}{N} \sum (x_i - \mu)^2$$

If we extract both μ and σ^2 from the same set of measurements:

$$\sigma^2 = \frac{1}{N} \sum (x_i - \bar{x})^2$$

ML variance estimator is biased! Bessel's correction missing - lecture 03.

Multiple parameter estimate

Likelihood function (and log-likelihood) can depend on multiple parameters:

$$\boldsymbol{\lambda} = (\lambda_1 \dots \lambda_p) \quad L = \prod_{j=1}^N f(\mathbf{x}^{(j)}; \boldsymbol{\lambda}) \quad \ell = \sum_{j=1}^N \ln f(\mathbf{x}^{(j)}; \boldsymbol{\lambda})$$

Multiple parameter estimate

Likelihood function (and log-likelihood) can depend on multiple parameters:

$$\boldsymbol{\lambda} = (\lambda_1 \dots \lambda_p) \quad L = \prod_{j=1}^N f(\mathbf{x}^{(j)}; \boldsymbol{\lambda}) \quad \ell = \sum_{j=1}^N \ln f(\mathbf{x}^{(j)}; \boldsymbol{\lambda})$$

Best estimate of $\boldsymbol{\lambda}$, for given set of experimental results $\mathbf{x}^{(j)}$, corresponds to maximum of the likelihood function, which can be found by solving a system of equations:

$$\left. \frac{\partial \ell}{\partial \lambda_i} \right|_{i=1 \dots p} = 0$$

Multiple parameter estimate

Likelihood function (and log-likelihood) can depend on multiple parameters:

$$\boldsymbol{\lambda} = (\lambda_1 \dots \lambda_p) \quad L = \prod_{j=1}^N f(\mathbf{x}^{(j)}; \boldsymbol{\lambda}) \quad \ell = \sum_{j=1}^N \ln f(\mathbf{x}^{(j)}; \boldsymbol{\lambda})$$

Best estimate of $\boldsymbol{\lambda}$, for given set of experimental results $\mathbf{x}^{(j)}$, corresponds to maximum of the likelihood function, which can be found by solving a system of equations:

$$\left. \frac{\partial \ell}{\partial \lambda_i} \right|_{i=1 \dots p} = 0$$

The Likelihood Principle

G. Bohm and G. Zech

Given a p.d.f. $f(\mathbf{x}; \boldsymbol{\lambda})$ containing an unknown parameters of interest $\boldsymbol{\lambda}$ and observations $\mathbf{x}^{(j)}$, **all information relevant for the estimation of the parameters $\boldsymbol{\lambda}$** is contained in the likelihood function $L(\boldsymbol{\lambda}; \mathbf{x}) = \prod f(\mathbf{x}^{(j)}; \boldsymbol{\lambda})$.

Multivariate Normal Distribution

Consider experiment resulting in a measurement $\mathbf{x} = (x_1, \dots, x_n)$.

If we assume each variable follows Gaussian p.d.f, the most general form of the joint probability distribution is:

$$f(\mathbf{x}; \boldsymbol{\lambda}) = A \exp \left[-\frac{1}{2} (\mathbf{x} - \boldsymbol{\lambda})^\top \mathbb{B} (\mathbf{x} - \boldsymbol{\lambda}) \right]$$

where $\boldsymbol{\lambda}$ is parameter vector and \mathbb{B} is an $n \times n$ matrix.

Since p.d.f. is symmetric about the point $\mathbf{x} = \boldsymbol{\lambda}$:

$$\begin{aligned} \mathbb{E}(\mathbf{x} - \boldsymbol{\lambda}) &= \int d\mathbf{x} (\mathbf{x} - \boldsymbol{\lambda}) f(\mathbf{x}; \boldsymbol{\lambda}) = 0 \\ \Rightarrow \mathbb{E}(\mathbf{x}) &= \boldsymbol{\lambda} \end{aligned}$$

Multivariate Normal Distribution

We should note that the derivative of the probability distribution:

$$\frac{\partial f}{\partial \lambda_i} = [(\mathbf{x} - \boldsymbol{\lambda})^\top \mathbb{B}]_i \cdot f(\mathbf{x}; \boldsymbol{\lambda})$$

Multivariate Normal Distribution

We should note that the derivative of the probability distribution:

$$\frac{\partial f}{\partial \lambda_i} = [(\mathbf{x} - \boldsymbol{\lambda})^\top \mathbb{B}]_i \cdot f(\mathbf{x}; \boldsymbol{\lambda})$$

We can now differentiate the formula for $\mathbb{E}(\mathbf{x} - \boldsymbol{\lambda})$ with respect to $\boldsymbol{\lambda}$:

$$\int d\mathbf{x} [(\mathbf{x} - \boldsymbol{\lambda}) (\mathbf{x} - \boldsymbol{\lambda})^\top \mathbb{B} - \mathbb{I}] \cdot f(\mathbf{x}; \boldsymbol{\lambda}) = 0$$

Multivariate Normal Distribution

We should note that the derivative of the probability distribution:

$$\frac{\partial f}{\partial \lambda_i} = [(\mathbf{x} - \boldsymbol{\lambda})^\top \mathbb{B}]_i \cdot f(\mathbf{x}; \boldsymbol{\lambda})$$

We can now differentiate the formula for $\mathbb{E}(\mathbf{x} - \boldsymbol{\lambda})$ with respect to $\boldsymbol{\lambda}$:

$$\int d\mathbf{x} [(\mathbf{x} - \boldsymbol{\lambda}) (\mathbf{x} - \boldsymbol{\lambda})^\top \mathbb{B} - \mathbb{I}] \cdot f(\mathbf{x}; \boldsymbol{\lambda}) = 0$$

and realizing that \mathbb{B} and \mathbb{I} are constant we get

$$\left[\int d\mathbf{x} (\mathbf{x} - \boldsymbol{\lambda}) (\mathbf{x} - \boldsymbol{\lambda})^\top \cdot f(\mathbf{x}; \boldsymbol{\lambda}) \right] \mathbb{B} = \mathbb{I}$$

Multivariate Normal Distribution

We should note that the derivative of the probability distribution:

$$\frac{\partial f}{\partial \lambda_i} = [(\mathbf{x} - \boldsymbol{\lambda})^\top \mathbb{B}]_i \cdot f(\mathbf{x}; \boldsymbol{\lambda})$$

We can now differentiate the formula for $\mathbb{E}(\mathbf{x} - \boldsymbol{\lambda})$ with respect to $\boldsymbol{\lambda}$:

$$\int d\mathbf{x} [(\mathbf{x} - \boldsymbol{\lambda}) (\mathbf{x} - \boldsymbol{\lambda})^\top \mathbb{B} - \mathbb{I}] \cdot f(\mathbf{x}; \boldsymbol{\lambda}) = 0$$

and realizing that \mathbb{B} and \mathbb{I} are constant we get

$$\left[\int d\mathbf{x} (\mathbf{x} - \boldsymbol{\lambda}) (\mathbf{x} - \boldsymbol{\lambda})^\top \cdot f(\mathbf{x}; \boldsymbol{\lambda}) \right] \mathbb{B} = \mathbb{I}$$

$$\Rightarrow \mathbb{C}_{\mathbf{x}} \mathbb{B} = \mathbb{I} \quad \Rightarrow \quad \mathbb{C}_{\mathbf{x}} = \mathbb{B}^{-1}$$

Multivariate Normal Distribution

We can now write the joint probability distribution as:

$$f(\mathbf{x}; \boldsymbol{\lambda}) = A \exp \left[-\frac{1}{2} (\mathbf{x} - \boldsymbol{\lambda})^T \mathbb{C}^{-1} (\mathbf{x} - \boldsymbol{\lambda}) \right]$$

where \mathbb{C} is the covariance matrix of variables \mathbf{x} .

Multivariate Normal Distribution

We can now write the joint probability distribution as:

$$f(\mathbf{x}; \boldsymbol{\lambda}) = A \exp \left[-\frac{1}{2} (\mathbf{x} - \boldsymbol{\lambda})^T \mathbb{C}^{-1} (\mathbf{x} - \boldsymbol{\lambda}) \right]$$

where \mathbb{C} is the covariance matrix of variables \mathbf{x} . **Log-likelihood:**

$$\ell(\boldsymbol{\lambda}; \mathbf{x}) = -\frac{1}{2} (\mathbf{x} - \boldsymbol{\lambda})^T \mathbb{C}^{-1} (\mathbf{x} - \boldsymbol{\lambda}) + \text{const}$$

$$\frac{\partial \ell}{\partial \lambda_i} = [(\mathbf{x} - \boldsymbol{\lambda})^T \mathbb{C}^{-1}]_i \quad \text{similar to p.d.f. derivative}$$

$$\frac{\partial^2 \ell}{\partial \lambda_i \partial \lambda_j} = - [\mathbb{C}^{-1}]_{ij} \quad \Rightarrow \quad \mathbb{C} = \left(-\frac{\partial^2 \ell}{\partial \lambda_i \partial \lambda_j} \right)^{-1}$$

Parameter covariance matrix

For the considered case of multivariate normal distribution, best parameter estimates $\hat{\lambda}$ are given by the measured variable values \mathbf{x} .

Unlike parameters λ , parameter estimates $\hat{\lambda}$ are random variables (functions of \mathbf{x}) and so we can consider covariance matrix for $\hat{\lambda}$:

$$\mathbb{C}_{\mathbf{x}} = \mathbb{C}_{\hat{\lambda}} = \left(-\frac{\partial^2 \ell}{\partial \lambda_i \partial \lambda_j} \right)^{-1}$$

Parameter covariance matrix

For the considered case of multivariate normal distribution, best parameter estimates $\hat{\lambda}$ are given by the measured variable values \mathbf{x} .

Unlike parameters λ , parameter estimates $\hat{\lambda}$ are random variables (functions of \mathbf{x}) and so we can consider covariance matrix for $\hat{\lambda}$:

$$\mathbb{C}_{\mathbf{x}} = \mathbb{C}_{\hat{\lambda}} = \left(-\frac{\partial^2 \ell}{\partial \lambda_i \partial \lambda_j} \right)^{-1}$$

Knowing the likelihood function, we can not only estimate parameter values, but also extract uncertainties and correlations of these estimates!

Parameter covariance matrix

For the considered case of multivariate normal distribution, best parameter estimates $\hat{\lambda}$ are given by the measured variable values \mathbf{x} .

Unlike parameters λ , parameter estimates $\hat{\lambda}$ are random variables (functions of \mathbf{x}) and so we can consider covariance matrix for $\hat{\lambda}$:

$$\mathbb{C}_{\mathbf{x}} = \mathbb{C}_{\hat{\lambda}} = \left(-\frac{\partial^2 \ell}{\partial \lambda_i \partial \lambda_j} \right)^{-1}$$

Knowing the likelihood function, we can not only estimate parameter values, but also extract uncertainties and correlations of these estimates!

For the uncorrelated parameters (diagonal covariance matrix):

$$\sigma_{\hat{\lambda}_i} = \left(-\frac{\partial^2 \ell}{\partial \lambda_i^2} \right)^{-1/2}$$

Parameter covariance matrix

Considered example was based on the Gaussian distribution.

Standard deviation is one of the parameters of the p.d.f., can be easily extracted from log-likelihood:

$$\sigma_i = \sqrt{C_{ii}}$$

However, this procedure works only in the Gaussian approximation.

How to define parameter uncertainty in the general case?

Parameter covariance matrix

Considered example was based on the Gaussian distribution.

Standard deviation is one of the parameters of the p.d.f., can be easily extracted from log-likelihood:

$$\sigma_i = \sqrt{C_{ii}}$$

However, this procedure works only in the Gaussian approximation.

How to define parameter uncertainty in the general case?

For 1-D case we have:

$$\ell(\hat{\lambda} + \sigma; \mathbf{x}) = \ln f(\mathbf{x}; \hat{\lambda} + \sigma) = \ell(\hat{\lambda}; \mathbf{x}) - \frac{1}{2}$$

Parameter covariance matrix

Considered example was based on the Gaussian distribution.

Standard deviation is one of the parameters of the p.d.f., can be easily extracted from log-likelihood:

$$\sigma_i = \sqrt{C_{ii}}$$

However, this procedure works only in the Gaussian approximation.

How to define parameter uncertainty in the general case?

For 1-D case we have:

$$\ell(\hat{\lambda} + 2\sigma; x) = \ln f(x; \hat{\lambda} + 2\sigma) = \ell(\hat{\lambda}; x) - \frac{4}{2}$$

Parameter covariance matrix

Considered example was based on the Gaussian distribution.

Standard deviation is one of the parameters of the p.d.f., can be easily extracted from log-likelihood:

$$\sigma_i = \sqrt{C_{ii}}$$

However, this procedure works only in the Gaussian approximation.

How to define parameter uncertainty in the general case?

For 1-D case we have:

$$\ell(\hat{\lambda} + 3\sigma; x) = \ln f(x; \hat{\lambda} + 3\sigma) = \ell(\hat{\lambda}; x) - \frac{9}{2}$$

Parameter covariance matrix

Considered example was based on the Gaussian distribution.

Standard deviation is one of the parameters of the p.d.f., can be easily extracted from log-likelihood:

$$\sigma_i = \sqrt{C_{ii}}$$

However, this procedure works only in the Gaussian approximation.

How to define parameter uncertainty in the general case?

Recipe for a parameter uncertainty

G. Bohm and G. Zech

Standard error intervals of the extracted parameter are defined by the decrease of the log-likelihood function by 0.5 for one, by 2 for two and by 4.5 for three standard deviations.

This definition works for arbitrary p.d.f. shape, also for multiple parameters

Multiple parameter estimate

Example log-likelihood function contours for a **sample of 10 events** from normal distribution, with extracted parameter values $\hat{\mu} = 1$ and $\hat{\sigma} = 2$.




Figure from: G. Bohm and G. Zech, Introduction to Statistics and Data Analysis for Physicists, Verlag Deutsches Elektronen-Synchrotron, 3rd edition

Multiple parameter estimate

For the set \mathbf{x} of N measurements we can write:

$$\begin{aligned}\sum (x_i - \mu)^2 &= \sum (x_i^2 - 2\mu x_i + \mu^2) \\ &= N (\langle x^2 \rangle - 2\mu \langle x \rangle + \mu^2)\end{aligned}$$

Multiple parameter estimate

For the set \mathbf{x} of N measurements we can write:

$$\begin{aligned}\sum (x_i - \mu)^2 &= \sum (x_i^2 - 2\mu x_i + \mu^2) \\ &= N (\langle x^2 \rangle - 2\mu \langle x \rangle + \mu^2) \\ &= N (\hat{\sigma}^2 + \hat{\mu}^2 - 2\mu \hat{\mu} + \mu^2) \quad \hat{\sigma}^2 = \langle x^2 \rangle - \langle x \rangle^2\end{aligned}$$

Multiple parameter estimate

For the set \mathbf{x} of N measurements we can write:

$$\begin{aligned}\sum (x_i - \mu)^2 &= \sum (x_i^2 - 2\mu x_i + \mu^2) \\ &= N (\langle x^2 \rangle - 2\mu \langle x \rangle + \mu^2) \\ &= N (\hat{\sigma}^2 + \hat{\mu}^2 - 2\mu \hat{\mu} + \mu^2) \quad \hat{\sigma}^2 = \langle x^2 \rangle - \langle x \rangle^2\end{aligned}$$

Log-likelihood function for the example is then:

$$\ell(\mu, \sigma; \hat{\mu}, \hat{\sigma}) = -\frac{N}{2\sigma^2} (\hat{\sigma}^2 + (\mu - \hat{\mu})^2) - \frac{N}{2} \ln(\sigma^2) + \text{const}$$


This corresponds to the Gaussian shape for μ , but very asymmetric for σ ...

Also, the two parameters are not independent!

Multiple parameter estimate

Example log-likelihood function contours for a **sample of 10 events** from normal distribution, with extracted parameter values $\hat{\mu} = 1$ and $\hat{\sigma} = 2$.

Log-likelihood contours for example measurement



Result from G. Bohm and G. Zech is nicely reproduced...

Multiple parameter estimate

As mentioned before, $\hat{\mu}$ and $\hat{\sigma}$ (extracted from measurements \mathbf{x}) are random variables! We can calculate their joint probability distribution:

$$f(\hat{\mu}, \hat{\sigma}^2; \mu, \sigma) = A \cdot \exp\left(-\frac{N(\hat{\mu} - \mu)^2}{2\sigma^2}\right) \cdot \left(\frac{N\hat{\sigma}^2}{\sigma^2}\right)^{k-1} \exp\left(-\lambda \frac{N\hat{\sigma}^2}{\sigma^2}\right)$$

where $\frac{N\hat{\sigma}^2}{\sigma^2}$ is distributed according to the Gamma distribution with $k = (N - 1)/2$ and $\lambda = 1/2$ (particular case referred to as χ^2 distribution; we will discuss it at the next lectures).

Multiple parameter estimate

As mentioned before, $\hat{\mu}$ and $\hat{\sigma}$ (extracted from measurements \mathbf{x}) are random variables! We can calculate their joint probability distribution:

$$f(\hat{\mu}, \hat{\sigma}^2; \mu, \sigma) = A \cdot \exp\left(-\frac{N(\hat{\mu} - \mu)^2}{2\sigma^2}\right) \cdot \left(\frac{N\hat{\sigma}^2}{\sigma^2}\right)^{k-1} \exp\left(-\lambda \frac{N\hat{\sigma}^2}{\sigma^2}\right)$$

where $\frac{N\hat{\sigma}^2}{\sigma^2}$ is distributed according to the Gamma distribution with $k = (N - 1)/2$ and $\lambda = 1/2$ (particular case referred to as χ^2 distribution; we will discuss it at the next lectures).

The two variables, $\hat{\mu}$ and $\hat{\sigma}$, are independent!

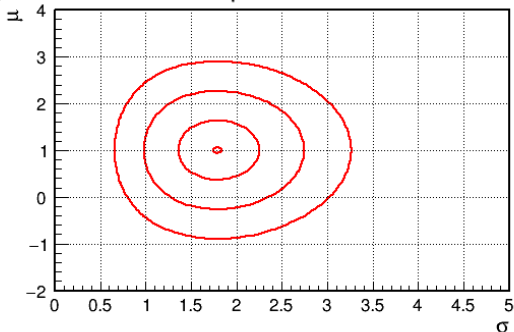
PDF for $\hat{\sigma}$ is asymmetric, but much less than the likelihood function !!!

Multiple parameter estimate

One needs to stress that likelihood function for p.d.f. parameters are not equivalent to probability distribution of parameter estimators!

Contours of the joint probability distribution function $f(\hat{\mu}, \hat{\sigma}; \mu, \sigma)$

Log-likelihood contours for example measurement




Contours corresponding to $\log f$ decrease by 0.5, 2 and 4.5 from maximum...

Multiple parameter estimate

One needs to stress that likelihood function for p.d.f. parameters are not equivalent to probability distribution of parameter estimators!

Results of Monte Carlo simulation (contours from 1 000 000 experiments)

Distribution of estimated mean vs sigma



In each experiment, mean and sigma are calculated from 10 generated numbers

Parameter Inference

- 1 Maximum Likelihood Method
- 2 Confidence intervals
- 3 Real life example

Presenting measurement results

When doing the measurement, we usually quote the final result as numerical value (with units) and estimated uncertainty:

$$x \pm \sigma_x$$

We can often calculate the uncertainty from the data itself (eg. when result is obtained by averaging a large number of independent measurements) or from the variation of the log-likelihood function.

Attributing proper uncertainty to the result is crucial!


But what does it tell us after all?!

Normal distribution

Meaning of σ is well defined for **Gaussian distribution**.

Probability for the experimental result to be consistent with the true value within $\pm N\sigma$ interval:

	α	
$\pm 1\sigma$	\Rightarrow 68.27	%
$\pm 2\sigma$	\Rightarrow 95.45	%
$\pm 3\sigma$	\Rightarrow 99.73	%
$\pm 4\sigma$	\Rightarrow 99.9937	%
$\pm 5\sigma$	\Rightarrow 99.999943	%




There is a non-zero chance for deviation greater than 5σ , but it is extremely small

Normal distribution

Meaning of σ is well defined for **Gaussian distribution**.

Probability for the experimental result to **differ from the true value** by more than $N\sigma$:

	α	
$\pm 1 \sigma$	$\Rightarrow 31.73$	%
$\pm 2 \sigma$	$\Rightarrow 4.55$	%
$\pm 3 \sigma$	$\Rightarrow 0.27$	%
$\pm 4 \sigma$	$\Rightarrow 0.0063$	%
$\pm 5 \sigma$	$\Rightarrow 0.000057$	%



Fluctuations up and down are observed with equal probability...

Normal distribution

Results of the Monte Carlo test ($\mu = 0$, $\sigma = 1$, 100 000 000 generations)

Down fluctuations:

$< -1 \sigma$ $p = 0.15864225$

$< -2 \sigma$ $p = 0.0227686$

$< -3 \sigma$ $p = 0.00134872$

$< -4 \sigma$ $p = 3.204E-05$

$< -5 \sigma$ $p = 3.2E-07$

Fluctuations up:

$> 1 \sigma$ $p = 0.15861607$


$> 2 \sigma$ $p = 0.02275479$

$> 3 \sigma$ $p = 0.00135162$

$> 4 \sigma$ $p = 3.263E-05$

$> 5 \sigma$ $p = 2.8E-07$

Probability distribution from MC



Good agreement with expectations

Poisson distribution

Results of the Monte Carlo test ($\mu = 10$, 100 000 000 generations)


Down fluctuations:

$< -1 \sigma$	$p = 0.1301325$
$< -2 \sigma$	$p = 0.01033324$
$< -3 \sigma$	$p = 4.55E-05$
$< -4 \sigma$	$p = 0$
$< -5 \sigma$	$p = 0$

Fluctuations up:

$> 1 \sigma$	$p = 0.13553469$
$> 2 \sigma$	$p = 0.02702754$
$> 3 \sigma$	$p = 0.00344483$
$> 4 \sigma$	$p = 0.00029364$
$> 5 \sigma$	$p = 1.774e-05$

Probability distribution from MC



Much longer tail of positive fluctuations!

Gamma distribution

Results of the Monte Carlo test ($\mu = \sigma^2 = 10$, 10 000 000 generations)

Down fluctuations:

$< -1 \sigma$ $p = 0.1533757$

$< -2 \sigma$ $p = 0.0046445$

$< -3 \sigma$ $p = 0$

$< -4 \sigma$ $p = 0$

$< -5 \sigma$ $p = 0$

Fluctuations up:

$> 1 \sigma$ $p = 0.1554444$


$> 2 \sigma$ $p = 0.0368358$

$> 3 \sigma$ $p = 0.0067406$

$> 4 \sigma$ $p = 0.0010105$

$> 5 \sigma$ $p = 0.0001305$

Probability distribution from MC




Even longer, 5σ fluctuations not excluded

Normal distribution in N-D

It is also important to notice that the fractions presented previously (eg. 68% within $\pm 1\sigma$) refer to one-dimensional normal distribution only!

If we consider 2-D distribution



Fractions within $N\sigma$ contours:

Deviation	Dimension			
	1	2	3	4
1 σ	0.683	0.393	0.199	0.090
2 σ	0.954	0.865	0.739	0.594
3 σ	0.997	0.989	0.971	0.939
4 σ	1.	1.	0.999	0.997

Less than 40% is contained inside 1σ contour...

G. Bohm and G. Zech

Interpreting results

As demonstrated above, quoting numerical result with uncertainties gives only partial information on the measurement...

It is sufficient, if we can assume normal distribution of the variable.

Also, we need to assume that the width of the distribution does not depend on the measured parameter. Only then the likelihood function will be Gaussian as well...

In the general case, parameter uncertainty does not give us full information on the tails of the distribution...

How should we present results of the experiment, if we are more interested in the probability of (large) result fluctuations?...

Interpreting results

There is also another problem, which has to be noticed!

So far we have only considered distribution of experimental results for given probability distribution, $f(\mathbf{x}; \boldsymbol{\lambda})$, when the parameter values $\boldsymbol{\lambda}$ are known.

Interpreting results

There is also another problem, which has to be noticed!

So far we have only considered distribution of experimental results for given probability distribution, $f(\mathbf{x}; \boldsymbol{\lambda})$, when the parameter values $\boldsymbol{\lambda}$ are known.

The actual situation is usually different: for given set of measurements \mathbf{x} we extract estimates of the parameter values $\hat{\boldsymbol{\lambda}}$.

Uncertainties estimated from log-likelihood variation indicate the expected level of agreement (in Gaussian approximation) between our estimate $\hat{\boldsymbol{\lambda}}$ and the true parameter values $\boldsymbol{\lambda}$.

Can we present measurement results in a way which gives us more precise information about the possible fluctuations?

Interpreting results

There is also another problem, which has to be noticed!

So far we have only considered distribution of experimental results for given probability distribution, $f(\mathbf{x}; \boldsymbol{\lambda})$, when the parameter values $\boldsymbol{\lambda}$ are known.

The actual situation is usually different: for given set of measurements \mathbf{x} we extract estimates of the parameter values $\hat{\boldsymbol{\lambda}}$.

Uncertainties estimated from log-likelihood variation indicate the expected level of agreement (in Gaussian approximation) between our estimate $\hat{\boldsymbol{\lambda}}$ and the true parameter values $\boldsymbol{\lambda}$.

Can we present measurement results in a way which gives us more precise information about the possible fluctuations?

Yes, but we need to define the problem differently...

We should not ask about the probability of $\hat{\boldsymbol{\lambda}}$...

Frequentist confidence intervals

Classical (frequentist) definition of the confidence interval refers directly to the probability distribution, $f(\mathbf{x}; \lambda)$.

For given outcome of the experiment x_m , $1 - \alpha$ confidence level (C.L.) interval for parameter λ is $[\lambda_1, \lambda_2]$, if for all values $\lambda' \in [\lambda_1, \lambda_2]$, our result x_m is inside the corresponding $1 - \alpha$ probability interval for $f(x; \lambda')$.

Frequentist confidence intervals

Classical (frequentist) definition of the confidence interval refers directly to the probability distribution, $f(\mathbf{x}; \lambda)$.

For given outcome of the experiment x_m , $1 - \alpha$ confidence level (C.L.) interval for parameter λ is $[\lambda_1, \lambda_2]$, if for all values $\lambda' \in [\lambda_1, \lambda_2]$, our result x_m is inside the corresponding $1 - \alpha$ probability interval for $f(x; \lambda')$.

This definition clearly depends on the way we define probability intervals for $f(x; \lambda')$ - it is rather a concept, more assumptions are needed.

We always refer to probability distribution for x !

Frequentist confidence intervals

It is interesting to note that the modern concept of confidence intervals was proposed by Polish statistician **Jerzy Neyman** only in 1937

Jerzy Neyman obtained his PhD (1924) and habilitation (1928) at UW

X—Outline of a Theory of Statistical Estimation Based on the Classical Theory of Probability

By J. NEYMAN

Reader in Statistics, University College, London

(Communicated by H. JEFFREYS, F.R.S.—Received 20 November, 1936—Read 17 June, 1937)

CONTENTS

	Page
I—INTRODUCTORY	333
(a) General Remarks, Notation, and Definitions	333
(b) Review of the Solutions of the Problem of Estimation Advanced Hereto	343
(c) Estimation by Unique Estimate and by Interval	346
II—CONFIDENCE INTERVALS	347
(a) Statement of the Problem	347
(b) Solution of the Problem of Confidence Intervals	350
(c) Example I	356
(d) Example II	362
(e) Family of Similar Regions Based on a Sufficient System of Statistics	364
(f) Example IIa	367

J.Neyman, *Phil. Trans. Royal Soc. London, Series A*, 236 333-80 (1937).

Frequentist confidence intervals

As mentioned above, to define confidence interval for parameter, we need to define how the probability interval for our measurement is defined.

There are three “natural” choices:

- We constrain the measurement from above:

$$\int_{x_{ul}}^{+\infty} dx f(x; \lambda) = \alpha$$

- We constrain the measurement from below:

$$\int_{-\infty}^{x_{ll}} dx f(x; \lambda) = \alpha$$

- We use central probability interval: **as presented for Gaussian pdf**

$$\int_{-\infty}^{x_1} dx f(x; \lambda) = \alpha/2 \quad \text{and} \quad \int_{x_2}^{+\infty} dx f(x; \lambda) = \alpha/2$$

Frequentist confidence intervals

Let us consider the simplest possible example, Gaussian pdf:

$$f(x; \lambda, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2} \frac{(x - \lambda)^2}{\sigma^2}\right)$$

assuming σ is known (and fixed).

What is the central 90% C.L. interval for λ , if we measure $x = x_m$?

Frequentist confidence intervals

Let us consider the simplest possible example, Gaussian pdf:

$$f(x; \lambda, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2} \frac{(x - \lambda)^2}{\sigma^2}\right)$$

assuming σ is known (and fixed).

What is the central 90% C.L. interval for λ , if we measure $x = x_m$?

From the Gaussian pdf properties we can directly obtain:

$$x_1 = \lambda - 1.64\sigma \quad x_2 = \lambda + 1.64\sigma$$

Frequentist confidence intervals

Let us consider the simplest possible example, Gaussian pdf:

$$f(x; \lambda, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2} \frac{(x - \lambda)^2}{\sigma^2}\right)$$

assuming σ is known (and fixed).

What is the central 90% C.L. interval for λ , if we measure $x = x_m$?

From the Gaussian pdf properties we can directly obtain:

$$x_1 = \lambda - 1.64\sigma \quad x_2 = \lambda + 1.64\sigma$$

Definition of the confidence interval for λ is based on the condition:

$$x_1 < x_m < x_2$$


Which is fulfilled for all λ in range:

$$\lambda_1 = x_m - 1.64\sigma < \lambda < x_m + 1.64\sigma = \lambda_2$$

So the central 95% C.L. interval for λ is $[x_m - 1.64\sigma, x_m + 1.64\sigma]$...

Frequentist confidence intervals


Graphical presentation of the procedure for $\sigma = 1$



G.J.Feldman, R.D.Cousins,
A Unified Approach to the Classical Statistical Analysis of Small Signals,
Phys.Rev.D57:3873-3889,1998; arXiv:physics/9711021

Frequentist confidence intervals


Graphical presentation of the general procedure



G.J.Feldman, R.D.Cousins,
A Unified Approach to the Classical Statistical Analysis of Small Signals,
Phys.Rev.D57:3873-3889,1998; arXiv:physics/9711021

Frequentist confidence intervals

General procedure




- calculate limits of probability intervals for x , $x_1(\theta)$ and $x_2(\theta)$, for different values of θ

R.L. Workman et al. (Particle Data Group), Prog. Theor. Exp. Phys. 2022, 083C01 (2022)
PDG web page

Frequentist confidence intervals

General procedure




- calculate limits of probability intervals for x , $x_1(\theta)$ and $x_2(\theta)$, for different values of θ
- calculated intervals define the “accepted region” in (θ, x)

R.L. Workman et al. (Particle Data Group), Prog. Theor. Exp. Phys. 2022, 083C01 (2022)
PDG web page

Frequentist confidence intervals

General procedure




- calculate limits of probability intervals for x , $x_1(\theta)$ and $x_2(\theta)$, for different values of θ
- calculated intervals define the “accepted region” in (θ, x)
- confidence interval for θ is defined by drawing line $x = x_m$ in the accepted region

R.L. Workman et al. (Particle Data Group), Prog. Theor. Exp. Phys. 2022, 083C01 (2022)
 PDG web page

Frequentist confidence intervals

General procedure

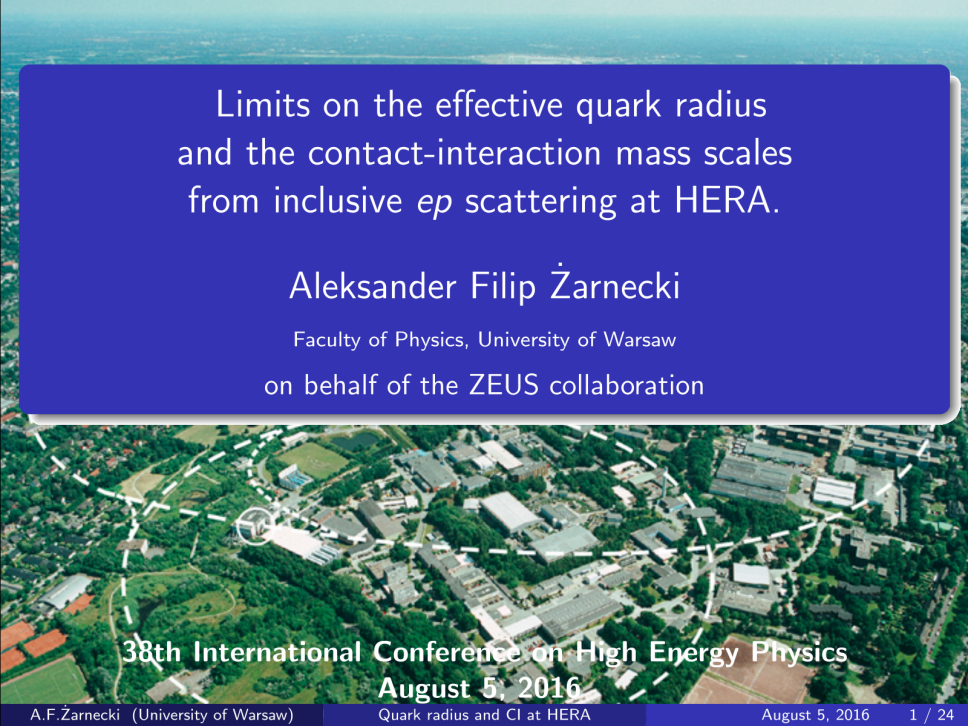


- calculate limits of probability intervals for x , $x_1(\theta)$ and $x_2(\theta)$, for different values of θ
 - calculated intervals define the “accepted region” in (θ, x)
 - confidence interval for θ is defined by drawing line $x = x_m$ in the accepted region
- \Rightarrow limit on θ for given x_m , $\theta_1(x_m)$, corresponds to limit on x for given θ : $x_m = x_1(\theta_1)$.

R.L. Workman et al. (Particle Data Group), Prog. Theor. Exp. Phys. 2022, 083C01 (2022)
 PDG web page

Parameter Inference

- 1 Maximum Likelihood Method
- 2 Confidence intervals
- 3 Real life example



Limits on the effective quark radius
and the contact-interaction mass scales
from inclusive ep scattering at HERA.

Aleksander Filip Żarnecki

Faculty of Physics, University of Warsaw
on behalf of the ZEUS collaboration

38th International Conference on High Energy Physics
August 5, 2016

Limits on contact interactions and leptoquarks at HERA

Aleksander Filip Żarnecki

Faculty of Physics, University of Warsaw
on behalf of the ZEUS collaboration



The European Physical Society Conference on High Energy Physics
Searches for New Physics parallel session

HERA

electron(positron)-proton collider at DESY



HERA I 1994–2000

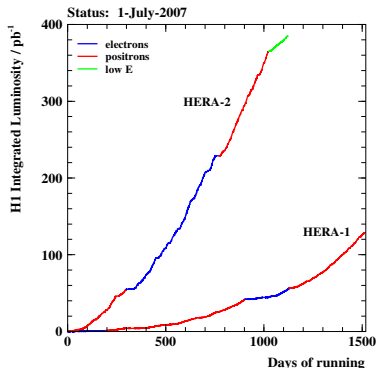
about 100pb^{-1} collected per experiment
 mainly e^+p data, unpolarised

HERA II 2002–2007

about 400pb^{-1} per experiment
 similar amount of e^-p and e^+p data

with longitudinal polarization of e^\pm beams (30–40%)


and small samples collected at reduced proton beam energy




Deep Inelastic $e^\pm p$ Scattering

Main process studied by H1 and ZEUS


NC DIS



CC DIS



Kinematic variables:



$$Q^2 = -(k - k')^2$$

|virtuality| of the exchanged boson

$$x = \frac{Q^2}{2P \cdot (k - k')}$$


fraction of proton momenta carried by struck quark

$$y = \frac{P \cdot (k - k')}{P \cdot k}$$

fraction of lepton energy transferred in the proton rest frame

SM predictions from HERA

H1 and ZEUS



NC and CC DIS cross sections comparable for the highest Q^2 values

$$Q^2 \sim M_Z^2, M_W^2$$

Combined QCD+EW analysis shows good agreement with SM predictions


Phys. Rev. D 93 (2016) 092002, [arXiv:1603.09628](https://arxiv.org/abs/1603.09628)

High precision data could also be used to look for possible BSM effects...

Quark form factor

“classical” method to look for possible fermion (sub)structure.

If a quark has finite size, the standard model cross-section is expected to decrease at high momentum transfer:



$$\frac{d\sigma}{dQ^2} = \frac{d\sigma^{SM}}{dQ^2} \cdot \left[1 - \frac{R_q^2}{6} Q^2\right]^2 \cdot \left[1 - \frac{R_e^2}{6} Q^2\right]^2$$

where R_q is the root mean-square radius of the electroweak charge distribution in the quark.

We do not consider the possibility of finite electron size...

same dependence expected for e^+p and e^-p !

QCD+BSM fit

Approach used for HERAPDF2.0 determination extended to take into account the possible BSM contribution

$$\chi^2(\mathbf{p}, \mathbf{s}, \eta) = \sum_i \frac{[m^i + \sum_j \gamma_j^i m^i s_j - \mu_0^i]^2}{(\delta_{i,\text{stat}}^2 + \delta_{i,\text{uncor}}^2) (\mu_0^i)^2} + \sum_j s_j^2$$

\mathbf{p} and \mathbf{s} are vectors of PDF parameters p_k and systematic shifts s_j ,
 η is the parameter describing BSM contribution (eg. $\eta = R_q^2$)

⇒ we fit them simultaneously to the combined HERA data

$$R_q^{\text{Data}} = -0.2 \cdot 10^{-33} \text{ cm}^2$$

μ_0^i and $m^i(\mathbf{p}, \eta)$ are measured and predicted (SM+BSM) cross sections,
 γ_j^i , $\delta_{i,\text{stat}}$ and $\delta_{i,\text{uncor}}$ are the relative correlated systematic, relative statistical and relative uncorrelated systematic uncertainties of the input data point i

Limit setting

Limits derived using the technique of MC replicas (**frequentist approach**).

Replicas are generated sets of cross-section values that are calculated for given $R_q^{2 \text{ True}}$ and varied randomly according to the statistical and systematic uncertainties (including correlations) of the input data.


Each replica is then used as an input to QCD+BSM fit

$$\Rightarrow R_q^{2 \text{ Fit}}$$

Number of replicas for each considered $R_q^{2 \text{ True}}$ value

\Rightarrow distribution of $R_q^{2 \text{ Fit}}$

$R_q^{2 \text{ True}}$ is tested by comparing $R_q^{2 \text{ Fit}}$ distribution with the value of $R_q^{2 \text{ Data}}$



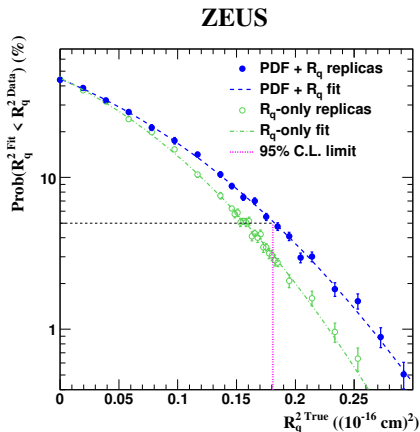
Limit setting

The probability of obtaining a $R_q^{2 \text{ Fit}}$ value smaller than that obtained for the actual data

$$\text{Prob}(R_q^{2 \text{ Fit}} < R_q^{2 \text{ Data}})$$

is studied as a function of $R_q^{2 \text{ True}}$


$R_q^{2 \text{ True}}$ values corresponding to the probability smaller than 5% are excluded at the 95% C.L.



$$R_q < 0.43 \cdot 10^{-16} \text{ cm}$$

limits obtained for fixed PDF parameters are too strong by about 10%








Results



$$-(0.47 \cdot 10^{-16} \text{ cm})^2 < R_q^2 < (0.43 \cdot 10^{-16} \text{ cm})^2$$

Framework

For many scenarios of “new physics” at much larger energy scale, BSM interactions can be approximated as $eeqq$ Contact Interactions (CI)



$eeqq$ contact interactions (CI)

Different CI scenarios assume different helicity structure of new interactions.

We consider different benchmark models with η as BSM parameter.

Limit setting

Limits derived using the technique of MC replicas (**frequentist approach**).

Replicas are generated sets of cross-section values that are calculated for given η^{True} and varied randomly according to the statistical and systematic uncertainties (including correlations) of the input data.


Each replica is then used as an input to QCD+BSM fit

$\Rightarrow \eta^{\text{Fit}}$

Number of replicas for each considered η^{True} value

\Rightarrow **distribution of η^{Fit}**

η^{True} is tested by comparing η^{Fit} **distribution** with the value of η^{Data}



Limit setting

Use replicas to calculate:


$$\text{Prob}(\eta^{\text{Fit}} < \eta^{\text{Data}})$$

for $\eta > \eta^{\text{Data}}$

$$\text{Prob}(\eta^{\text{Fit}} > \eta^{\text{Data}})$$

for $\eta < \eta^{\text{Data}}$

for different η^{True}



Limit setting

Use replicas to calculate:

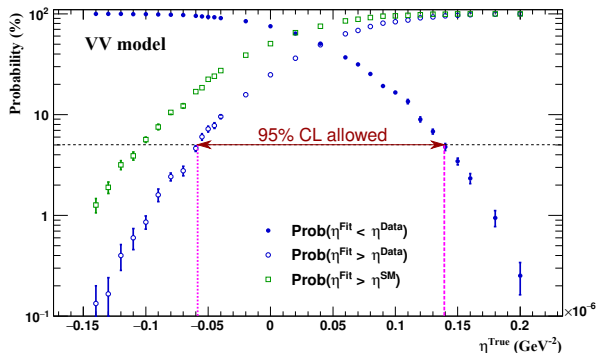
$$\text{Prob}(\eta^{\text{Fit}} < \eta^{\text{Data}})$$

for $\eta > \eta^{\text{Data}}$

$$\text{Prob}(\eta^{\text{Fit}} > \eta^{\text{Data}})$$

for $\eta < \eta^{\text{Data}}$


for different η^{True}



Excluded on 95% C.L. are η^{True} resulting in probability below 5%.

The limit-calculation procedure was repeated for systematic variations considered. The weakest of the obtained coupling limits was taken as the result of the analysis and used to calculate the final mass-scale limits.

Contact Interaction limits



limits calculated without and with modeling
uncertainties compared with the expected ones

ZEUS

HERA $e^{\pm}p$ 1994–2007 95% C.L. limits (TeV)

Model	Observed (exp+mod)		Expected		p_{SM} (%)
	Λ^-	Λ^+	Λ^-	Λ^+	
LL	12.8	4.5	5.9	6.3	7.0
RR	14.7	4.4	5.7	6.1	5.9
LR	4.7	5.5	5.7	6.3	34
RL	5.0	5.3	5.6	6.5	42
VV	13.9	9.0	11.2	11.4	25
AA	15.7	4.2	7.9	7.8	0.6
VA	3.6	3.5	4.2	4.2	5.8
X1		3.2	5.4	5.5	0.4
X2	10.4	6.4	7.8	8.3	24
X3	17.9	6.2	8.3	8.7	7.3
X4	7.2	7.5	8.0	8.6	39
X5	9.5	6.4	7.7	7.7	27
X6	3.1		5.3	5.5	0.3


Heavy Leptoquark limits

ZEUS

λ_{LQ}/M_{LQ} 95% C.L. limits (TeV^{-1})

ZEUS





HERA e^+p 1994–2007 95% C.L. Limits



Model	Observed (exp+mod)	Expected	p_{SM}
S_0^L	0.28	0.56	9.0
S_0^R	1.03	0.72	5.5
\tilde{S}_0^R		1.71	1.8
$S_{1/2}^L$	0.83	0.76	43
$S_{1/2}^R$	1.04	0.92	39
$\tilde{S}_{1/2}^L$	1.66	1.39	38
S_1^L	1.18	0.62	< 0.01
V_0^L		0.44	0.5
V_0^R	1.47	0.99	1.8
\tilde{V}_0^R	0.18	0.53	5.5
$V_{1/2}^L$	1.19	1.29	38
$V_{1/2}^R$	0.67	0.57	39
$\tilde{V}_{1/2}^L$	0.59	0.49	43
V_1^L	0.41	0.25	32

QCD+CI fit results

Improved description of the data for four models (3CI+1LQ): $\Delta\chi^2 < -4$.



X6: change dominated by CI contribution
 $\Delta\chi^2 = -6.01$

S_1^L : significant change in proton PDF !
 $\Delta\chi^2 = -11.1$