

Topical Lectures Statistics – Exercises Set 1

Wouter Verkerke (Dec 2022)

General Instructions

Input files

Input files for exercises can be found in three places

1. At nikhef (stbc-i1.nikhef.nl) in directory `~verkerke/stats2022/`
2. At CERN (lxplus7.cern.ch) in directory `~verkerke/public/stats2022`
3. On the web at <http://www.nikhef.nl/~verkerke/stats2022>

Running ROOT

All exercises are based on ROOT. You are recommended to use version 6.24.08,

in which all prepared material has been tested. To pick up a pre-installed ROOT

version please execute the following setup script

```
source /cvmfs/sft.cern.ch/lcg/app/releases/ROOT/6.24.08/x86_64-  
centos7-gcc48-opt/bin/thisroot.sh
```

This release will work both at Nikhef and at CERN

Where to work

If you have an account at Nikhef, please work on stbc-i1/2/3.nikhef.nl only and not on login.nikhef.nl (these run CentOS7 – required for above ROOT version)

If you have an account at CERN, you can also work on lxplus7.cern.ch, this will select an CentOS7 node (required for above ROOT version)

You can also work directly on your laptop if you have ROOT installed yourself

Quick reference guide to RooFit model building

In these tutorial exercises we will use RooFit to build probability models as that allows to rapidly build complex models. This requires some familiarity with the RooFit model building syntax. It is easy enough to pick up 'on the way' from input files given with this tutorial, but for completeness this page also presents a quick reference guide to the model building syntax and strategy

RooFit model structure

The key feature of RooFit model building is that all elements of a probability model (functions, p.d.f.s, variables) are all individual objects that connect to each other. For example a Gaussian probability density model is expressed as

```
RooRealVar x("x","x",0,-10,10) ;
RooRealVar mean("mean","mean",0,-10,10) ;
RooRealVar sigma("sigma","sigma",3,0.1,10) ;
RooGaussian gauss("gauss","gauss",x,mean,sigma) ;
```

The simplest strategy to build these models is to have all objects contain in a workspace, and use the 'factory' tool to fill the workspace with objects:

```
RooWorkspace w("w") ;
w.factory("Gaussian::gauss(x[0,-10,10],mean[0,-10,10],sigma[3,0.1,10])") ;
w.Print() ; // see what's inside the workspace
```

Inside the workspace this factory specification builds exactly the model shown above. The elements of the model can be accessed as follows:

```
RooAbsPdf* gauss = w.pdf("gauss") ; // extract a pdf
RooRealVar* x = w.var("x") ; // extract a variable
```

RooFit factory syntax

The factory syntax is quick to learn since it has very few rules

- To create a variable use `x[val,min,max]`.
- To create a pdf or function use `ClassName::objectName(args...)`
Any RooFit function or pdf class can be created this way.
You are allowed to omit the 'Roo' prefix from any class name.
The meaning and order of the args match those specified in the constructor of the class (after the mandatory name and title). Example:

```
C++ constructor: RooGaussian gauss("gauss","gauss",x,mean,sigma)
Factory spec: w.factory("Gaussian::gauss(x,mean,sigma)
```

- Any field where an object name is expected must either contain the name of a previously created object, or it can be created on this spot.

```
w.factory("x[0,-10,10]") ;
w.factory("Gaussian::gauss(x,mean[0,-10,10],sigma[3,0.1,10])") ;
```

In 2nd line above, for x, a previously created object is referenced, for mean and sigma, new objects are created in place.

Special Factory rules for operators

For certain functions and pdf a custom syntax exists to simplify their use. The most important ones are shown here

- **Addition of two (or more) pdfs**

```
SUM::objName(frac1*pdf1, pdf2) ; // shape-only pdf
SUM::objName(yield1*pdf1, yield2*pdf2) ; // extended pdf: N(exp) = y1+y2
```

- **Multiplication of two (or more) pdfs**

```
PROD::objName(pdf1, pdf2, ...) ; // product of independent pdfs
PROD::objName(pdf1|x, pdf2, ...) // product involving a conditional pdf
```

- **Interpreted function expressions** (all TFormula expressions can be used). All symbols reference in the expression must be passed as arguments.

```
expr::funcName('some expression', <list of objects used>) ;
```

Example use cases of these operators are given in the input files of various exercises and will further clarify their use. More special operator syntax exists, but is not needed for this course (e.g. for convolutions, joint models, integrals, projections, amplitude sums)

Basic use of RooFit models

Probability models in RooFit can be used for event generation, (likelihood) fitting and plotting. All of these are one-line operations.

- Generation of unbinned toy dataset

```
RooDataSet* pdf.generate(observables, eventcount) ;
RooDataSet* pdf.generate(observables) ; for extended models only
```

- Unbinned Likelihood fit of model to data

```
pdf.fitTo(data) ;
RooFitResult* r = pdf.fitTo(data, Save()) ; // save extra info
```

- Plotting of data and model

```
RooPlot* frame = obs.frame() ; // creates empty plot frame
data.plotOn(frame) ; // plot data on frame
pdf.plotOn(frame) ; // project pdf on obs, normalize to data
```

Questions? Don't hesitate to ask!

Exercise 1 – The Central Limit Theorem

In the lectures of Day 1 we saw that the Central Limit Theorem predicts that the sum of N measurements has a Gaussian distribution in the limit of $N \rightarrow \infty$, independent of the distribution of each individual measurement

- In this exercise we will investigate how quickly this convergence happens as function of N .
- We start with a ‘fake’ measurement resulting in a value x with a uniform distribution between $[0,1]$ (i.e., this is very non-Gaussian)
- Then we will look at the distribution of x_1+x_2 , $x_1+x_2+x_3$, etc and compare these with the properties of a Gaussian distribution

Start with file `ex08.C`

- This macro books a ROOT histogram, runs 10000 experiments and fills the ‘measured’ value of x in the histogram and plots the histogram and the end of the run.
- Look at the macro and run it (`root -l ex01.C` from the OS command line, or `.x ex01.C` from the ROOT command line)
- Run the macro again now passing value 2 as argument for `Nsum` `.x ex01.C(2)` (or `root -l 'ex01.C(2)'` from the OS command line. Note that in this case the quotations are essential). Look at the distribution. Repeat for `Nsum=3,5,10,20` and 100.

You will see that around `Nsum=10` the distribution is already looks quite Gaussian.

- Next, remove the break in the code by commenting the ‘return’ statement and run again. The code will draw a Gaussian distribution on top of the sampled distribution
- Math exercise: calculate the expected RMS of the distribution analytically, and insert that expression for the width of the Gaussian. Hint: first calculate the RMS of a uniform distribution $[0,1]$ ** , then calculate from that the variance of the sum of N distributions using the Central Limit Theorem.

** $\text{RMS} = \sqrt{V}$, with variance $V = E[x^2] - E[x]^2$, where $E[x]$ is the expectation value of x according to the sampled distribution (=uniform in range $[0,1]$)

- Observe how the distributions deviates from a Gaussian for `Nsum=3,10,100` (you may need to increase the statistics from the default 10K to much more for `Nsum=10,100` to be able to see the deviations well)

Exercise 2 – An unbinned maximum likelihood fit

This is a mostly a demonstration exercise only that shows some of the functionality and the syntax of the RooFit toolkit for data modeling that we'll be using in the next exercises

Copy file `ex02.c`, look at it and run it. This macro does the following

- It creates a Gaussian probability density function
- It generates an unbinned dataset with 10k events from that function
- It performs an unbinned ML fit of the Gaussian model to the toy dataset
- It makes a plot of the data and overlays it with the Gaussian model

Now comment out the part of the code labeled as 'block 1' and run the macro again

- This code will print out the covariance matrix and correlation matrix of the fit parameters.

Now comment out the code labeled 'block 2' and run again.

- This code will visualize the uncertainty of the model on the canvas using the error propagation technique. At 10K the event the uncertainty is very small (you can see it if you zoom in on the peak region of the pdf)
- Change the number of generated events from 10K to 100 and change the binning of the data in the plot from 100 bins to 10 bins (this is the argument in the `w.var("x")->frame()` call. Run again.
- Lower the number of generated events from 100 to 10 and run again. The error on the shape will now be significant, but you see that an unbinned ML can reliably fit very small event samples

Now comment out code block 3

- This will visualize the error on the p.d.f shape due to the uncertainty on the mean parameter only

Exercise 3 – Using the NP Lemma for event selection

The Neyman-Pearson lemma states that the optimal selection of selection of data points (x,y,z) that contain a mix background (hypothesis H_0) and signal (hypothesis H_1) is given by a contour of the likelihood ratio

$$L(H_0)/L(H_1) > c$$

where the critical value c can be chosen to obtain a selection with the desired tradeoff between purity (ratio of H_0 and H_1 in selection) and efficiency (fraction of H_2 in selection). We are going to exploit this NP lemma to create a plot projection of a three-dimensional pdf on its observable x , while applying an NP-optimal selection criteria on its observables y,z .

Copy file `ex03.c`. This macro performs the following steps

- Construct a 3D Gaussian signal model, and a 3D Polynomial bkg model
- Generate a sample of 10.000 events from the model
- Fit the model to the data
- Plot a plain projection of the data and model on observable x

As is clear from the plot, when all data is projected in a single observable, there substantial background under the signal, as the projection on x does not take the discriminating information in observables y,z into account.

- Uncomment the return statement half-way the macro. The macro will now additionally do the following
 - Construct a NP-optimal selection criteria $f L(H_0)/L(H_1)$ where H_0 =signal, and H_1 =signal+background, for events using the (y,z) information of the model only
 - Make a plot of the data projected on x , only selecting those events that meet $L(H_0)/L(H_1) > 0.7$
 - Make a corresponding projection of the 3-model that integrates out only region selected by the requirement $L(H_0)/L(H_1) > 0.7$
- Observe the much-improved signal-to-background ratio in the new plot
- Try to change the Likelihood ratio cut value up and down and see what happens? *Why does $L(H_0)/L(H_1) > 0.8$ result in a empty plot?*

Finally, we compare the sensitivity to signal significance for two potential workflows

- 1 – No event selection – big fit too all events
- 2 – NP-optimal event selection using observables y,z .
Subsequent fit on selected events in observable x only

The first workflow is already executed in the macro – it is simply the initial fit. The second workflow can be enabled by removing the return statement close to the end of the macro. *How do does the signal significance differ?* (In this regime of signal yields, you can approximate the significance as $Z=f_{\text{sig}}/\sigma(f_{\text{sig}})$). Repeat for several different settings of the Likelihood Ratio cut.

Exercise 4 – A simple two-dimensional model

A simple extension of the 1-dimensional signal+background model of exercise 4 involves a variant where the mean of the Gaussian signal changes as function of a second observable y in the data

Copy file `ex04.c`. This macro performs the following steps

- Construct an extended model that is the sum of Gaussian signal and an Exponential background
- Generate a toy dataset with 1000 events
- Perform an *unbinned* ML fit to the data
- Plot data and projection of the model on the observable x

Run the macro and observe its behavior.

- Do you understand the shape of the signal in the plot?
- Change the slope in the `mean_func()` of the Gaussian signal model to a smaller number (or even zero) and see how the shape changes
- Does the precision of the estimate of `Nsig` depend on the magnitude of the slope of `mean_func`? Explain why (not)?
- Do you understand the (anti)correlation between the estimates of `Nsig` and `Nbkg`?
- You can (as an option) also visualize the model in 2 dimensions by uncommenting the last part of the macro. You may need to rotate the plot from its default viewing angle by drag-clicking on it.

Exercise 5 – A sig+bkg fit & effects at low statistics

The most popular statistical models used in particle physics describe the data as a sum of a signal and background component in some discriminating observables. In this exercise we will explore the features of such a model, especially at very low sample sizes.

Copy file `ex05.c`. This macro performs the following steps

- Construct a model that is the sum of Gaussian signal and an Exponential background
- Generate a toy dataset with 1000 events
- Perform an *unbinned* ML fit to the data
- Plot data and projection of the model

Run the macro and observe its behavior. Then try a few variants

- Reduce the sample size to 100 and 10 and try again. If needed you can adjust the binning for the data in the plot projection (this will not affect the fit since it uses the data unbinned)
- Uncomment the bottom section that runs a MC Study of the properties of this model: it generates 1000 toy datasets, fits each of them and collects the fit statistics, by default on the `fsig` parameter. *Does the fit look unbiased?*
- Lower the sample size for the toys in the study from 1000 to 50, and then to 10 and run for each reduced size again. *Have your conclusions changed?*

One of the problems the fit runs into at low statistics is that the signal fraction `fsig` has a lower boundary at zero. If statistics are sufficiently small, ML fits will frequently run into this boundary, warping distributions.

- To mitigate boundary effects, decrease the lower bound on `fsig` from 0 to -1. Rerun the toy study for sample size 10. Does it improve? Rerun again for 100 to see how the behavior improves.

Finally, we push the model to see how well it can constrain `fsig` for data that doesn't contain any signal (*this strategy relates to understanding the potential for discovery of a signal – this will be discussed further in tomorrow's lectures*).

- Change the default `fsig` value to zero, and run the study for sample size 100. Then decrease the sample size to 10 and run again.

As you, at sufficiently low sample size, ML estimators will suffer from biases. ML estimators are unbiased, per the ML theorem, *if such an unbiased estimator exists*, but for very low statistics samples such estimators will not exist. Generally, bias terms scale with $1/N$, hence decreasing rapidly rising N also in comparison with the statistical error that scales with $1/\sqrt{N}$

Exercise 6 – Maximum Likelihood life time estimator

For certain simple cases it is possible to calculate the ML estimate of a parameter analytically rather than relying on a numeric estimate. A well-known case is that of the fit of a lifetime of an exponential distribution

Copy `ex06.c` and run it. This example performs an unbinned MLE fit to an exponential distribution of 100 events to estimate the lifetime.

Now we aim to construct the *analytical* ML estimator of the lifetime (do this part on paper, not by computer)

- Write down the probability density function for the exponential lifetime distribution.
- It is essential that you formulate a *normalized* expression, i.e. $\int_0^\infty F(t) dt = 1$ when integrated from 0 to ∞ .
The easiest way to accomplish that is to divide whatever you expression you have by the integral of that expression over dt . (*You can calculate that normalization integral on paper*)
- Next write down the analytical form of the negative log-likelihood $-\log(L)$ for that probability density function given a dataset of N events labeled these x_i in your expression. *Be sure to also include the pdf normalization term in the expression*
- The analytical ML estimate of the lifetime τ then follows the requirement that $d(-\log L)/d\tau = 0$. Calculate the expression for this derivative solve and derive the value of τ for which the requirement $d(-\log L)/d\tau$ holds.

Finally, implement the analytical calculation of MLE estimator for τ in the code of `ex06.c`

- Uncomment block one, which implements a look over the dataset, retrieving the values of the decay time one by one and build your calculation of the analytical estimate of τ with that of the numeric calculation from `fitTo()`
- Explain why you might have minor discrepancies between the analytical and numeric calculations.
- Increase the event count from 100 to 10000 and run again

Exercise 7 – A likelihood fit for an efficiency curve

This exercise is mostly a demonstration. The goal of this exercise is to demonstrate that every modeling problem can be cast in the form of a probability model, so that any subsequent parameter estimation can be done with a likelihood fit.

Copy `ex07.C`. This macro presents a classical problem that is usually seen as ‘ χ^2 -fit’ problem.

- A dataset has two observables: a continuous observable x , e.g., an energy of a triggered object, a label c that indicates if the event is accepted or rejected by the trigger.
- *The goal is to fit a turn-on curve of the efficiency of the trigger.* This efficiency function (which is clearly not a probability model itself!) has an S-curve in it that can be fit to describe a turn-on point in the data at some energy, and with some resolution (i.e. it is not a very sharp turn-on)
- The classical (‘naïve’) approach is to first construct an efficiency histogram from the data (fraction-accepted vs E). Then a χ^2 -fit is performed of the efficiency function $f(E|p)$ with some parameters p to the observed efficiency histogram.
- Run the macro – and see the ‘naïve’ approach executed as a plain ROOT χ^2 -fit

Apart from the fact that the χ^2 -fit must be a binned fit, the highly non-Gaussian nature of the binomial errors on the efficiency data (especially close to 0 or 1) is cause for concern. It would thus be preferable to not do this. How can this parameter estimation problem be captured in a likelihood model?

- The first step is to realize that one needs to go back to the original 2-dimensional dataset $D(E,c)$ where E is the energy observable and c is discrete labeling observable that can take states ‘accept’ and ‘reject’
- Given an efficiency $\epsilon(E)$ that is valid at energy E , the probability model for the observable c is a Binomial model with $n=2$.
- The conditional probability model for this data is thus simply

$$f(c|E) = \text{Binomial}(n=2, k=\epsilon(E))$$

- The full pdf would then be $f(c|E)*F(E)$, but this step can be omitted if there is no signal and background component to the full model as $F(E)$ could simply be taken to be identical the observed distribution of E .

Remove the return statement in the code and run the likelihood fit alongside the ROOT χ^2 -fit.

- Observe that for 1000 events both fits are comparable in precision
- Lower the event count to 100 events, or even 50 events, rerun and observe the superior stability of the unbinned likelihood-based efficiency fit.