

LHCb Starterkit 2022 - Practice Session

Part 2 - Analysing your data

2nd December 2022

Goals of the session

We will now put to use some of the knowledge learned in the first two days of this workshop. You (should) have the data you're interested in, in some ROOT ntuples. What can you do with it?

We will go through some basic steps that you would take to study a Monte Carlo sample of your signal channel of interest. This includes applying some rudimentary cuts to it, and fitting the resulting mass distribution.

1 Setup of the Session

You should have worked this morning on producing a ntuple for the decay $D^0 \rightarrow K^-\pi^+$, with a simple DaVinci script. As running DaVinci can take a long time to get a reasonable amount of signal events in your ntuple, for this second part of the Session you are provided with a ROOT file that contains the TTree with the kinematic, PID, and MC-truth variables, but with much higher statistics. This file can be copied on lxplus from

```
/afs/cern.ch/user/l/lopaoluc/public/PracticeSession/ntuple.root
```

We will be running *zfit* to fit the simulated data contained in the ntuple, so you will need to setup the correct python environment. As we have seen in the first two days of the session, in the main Starterkit folder that you have cloned from git ([here](#)), you should be able to set up the environment by just running on your command line:

```
source starterkitEnv/run bash
```

to open a bash shell in the correct environment.

You can write a python script or a jupyter notebook if you want to add comments and explanation for the steps that you'll take throughout the session. We will not mark you, so this is just for your own future reference!

Use the lessons and materials provided during the week as you go along! This Session is not meant to be test of your knowledge, but just to make you more familiar with the tools we've seen. Helpers will be around to assist with issues you may encounter.

2 Exploring the data

First, you can use `uproot` to read and load in the data contained in the ntuple. The TTree in the ROOT file is under `D0_KmPip/DecayTree`, so that is the path that you should use when loading it. You can start by loading in the D^0 mass branch, which is called `D0_M`. Try and plot it and see what it looks like! You should see a peak over a continuous background.

Later, you can start loading in more branches. Here's a list of the ones that will be useful for today:

```
D0_M, D0_TRUEID, Kminus_PIDK, Kminus_PIDpi, Kminus_PIDmu, Kminus_TRUEID,
piplus_PIDK, piplus_PIDpi, piplus_PIDmu, piplus_TRUEID
```

You can have a look at them by plotting them and seeing what their distributions look like. We can note a few things about these branches:

- The ntuple contains both $D^0 \rightarrow K^- \pi^+$ and the charge-conjugate process $\bar{D}^0 \rightarrow K^+ \pi^-$. The names of the branches of the children K^- and π^+ are only meant to reflect what the process is, but it doesn't mean that those branches *only* contains kaons and pions of a given charge.
- The `TRUEID` variables return an identification number assigned at MC level to any particle candidate in the ntuple. This is useful because simulated data that passes through the nominal LHCb reconstruction may produce candidates which do not correspond to the original signal generated by Gauss (as reconstruction is not 100% efficient). The `TRUEID` for D^0 , K^- and π^+ are 421, -321, 211, respectively. More information on the conventions for this numbering scheme can be found [here](#).
- The `PID` variables can be interpreted as a likelihood that a particle candidate is of a given species. For example, `Kminus_PIDK` returns the likelihood that the K^- candidate is actually a kaon; `Kminus_PIDmu` returns the likelihood that the K^- candidate is actually a muon; and so on. Note that the number returned is not a probability, *i.e.* a number from 0 to 1, but as a rule of thumb the chance of a particle being of a given species is higher the higher the value of the corresponding `PID`.

We can use these to apply cuts on our data! For example, when using simulation, a common selection is what is known as *truth-matching*. This means requiring that the `TRUEID` of the candidates match what we expect. Because we have both D^0 and \bar{D}^0 decay candidates, we have to require that the *absolute value* of the `TRUEID` matches the known values.

Apply the truth-matching cut on your data and look at how the shape of the distribution changes. You should see that most of the background gets removed, and only the peak remains.

Alternatively, you can play with the PID variables and apply different cuts to see how the distribution is affected.

3 Fitting your data

We're ready now to start doing some fits! We want to fit the mass distribution of the D^0 candidates. First, start by setting up your observable which will be the proxy for the variable `D0_M`.

As you have seen by plotting the data, we can start by modelling the peak as a simple gaussian. Set up the parameters for the mean and width, as well as a yield parameter. Try doing an extended unbinned fit with the gaussian to the non-truth matched data, and print out the result. Is it a good fit?

Next, you can plot the data and the fitted model together to see how well the gaussian models the data. You can also try to plot the 1D likelihood function for the mean of the gaussian.

Advanced: you can also try to plot what are known as *pulls*. You can think of pulls as a very naive measure of the goodness-of-fit. If you plot your data as a histogram, then the pull relative to one of its bins is defined as such:

$$pull = \frac{N_{bin} - y_{bin}}{\sqrt{N_{bin}}} \quad y_{bin} = \int_{bin} pdf(m) dm$$

This is the difference between the number of entries in a given bin N_{bin} with the integral of the extended pdf over that bin, normalised by the poissonian error $\sqrt{N_{bin}}$. Generally, in a good fit, the distribution of pulls for all bin should be centered in zero and be symmetric. They should be randomly scattered and no systematic trend should be visible.

4 Improving the model

Since we're not modelling the background at the moment, we shouldn't expect to see a good fit when plotting it against our data. So, we can add to our model a second pdf, which represents the background. This can be modelled by an exponential function. Remember to set up the yield for this as well and sum the resulting extended pdf with the one that we made for the gaussian peak.

Try to fit this updated model to the data, and plot it. Does it look better?

Advanced: Plot the pulls for this as well.

5 Advanced: Storing the result

Now comes the part for you to reflect and be creative. We want ideally to store the information of the fit somewhere, how can we do that?

Hint: Would dictionaries be useful in this case?