

*s*Plot : A QUICK INTRODUCTION

M. PIVK

CERN, 1211 Geneva 23, Switzerland

E-mail: mariel.pivk@cern.ch

The paper advocates the use of a statistical tool dedicated to the exploration of data samples populated by several sources of events. This new technique, called *s*Plot, is able to unfold the contributions of the different sources to the distribution of a data sample in a given variable. The *s*Plot tool applies in the context of a Likelihood fit which is performed on the data sample to determine the yields of the various sources.

1 Introduction

This paper describes a new technique to explore a data sample when the latter consists of several sources of events merged into a single sample of events. The events are assumed to be characterized by a set of variables which can be split into two components. The first component is a set of variables for which the distributions of all the sources of events are known: below, these variables are referred to as the *discriminating* variable. The second component is a set of variables for which the distributions of some sources of events are either truly unknown or considered as such: below, these variables are referred to as the *control* variables.

The new technique, termed *s*Plot^a, allows one to reconstruct the distributions for the control variable, independently for each of the various sources of events, without making use of any *a priori* knowledge on this variable. The aim is thus to use the knowledge available for the discriminating variables to be able to infer the behavior of the individual sources of events with respect to the control variable. An essential assumption for the *s*Plot technique to apply is that the control variable is uncorrelated with the discriminating variables.

The *s*Plot technique is developed in the context of a maximum Likelihood method making use of the discriminating variables. Section 2 is dedicated to the definition of fundamental objects necessary for the following. Section 3 presents an intermediate technique, simpler but inadequate, which is a first step towards the *s*Plot technique. The *s*Plot formalism is then developed Section 4 and its properties explained in Section 5. An example of *s*Plot at work is provided in Section 6 and some applications are

described in Section 7. Finally, the case where the control variable is correlated with the discriminating ones is discussed in Section 8.

2 Basics and definitions

One considers an unbinned extended maximum Likelihood analysis of a data sample in which are merged several species (signal and background) of events. The log-Likelihood is expressed as:

$$\mathcal{L} = \sum_{e=1}^N \ln \left\{ \sum_{i=1}^{N_s} N_i f_i(y_e) \right\} - \sum_{i=1}^{N_s} N_i, \quad (1)$$

where

- N is the total number of events considered,
- N_s is the number of species of events populating the data sample,
- N_i is the (non-integral) number of events expected on the average for the i^{th} species,
- y represents the set of discriminating variables, which can be correlated with each other,
- $f_i(y_e)$ is the value of the Probability Density Function (pdf) of y for the i^{th} species and for event e .

The log-Likelihood \mathcal{L} is a function of the N_s yields N_i and, possibly, of implicit free parameters designed to tune the pdfs on the data sample. These parameters as well as the yields N_i are determined by maximizing the above log-Likelihood.

The crucial point for the reliability of such an analysis is to use an exhaustive list of sources of events combined with an accurate description of all the pdfs f_i . If the distributions of the control variables are known (resp. unknown) for a particular

^aThe *s*Plot technique is the subject of a publication¹ where details of the calculations and more examples can be found.

source of events, one would like to compare the expected distribution for this source to the one extracted from the data sample (resp. determine the distribution for this source)^b.

The control variable x which, by definition, does not explicitly appear in the expression of \mathcal{L} , can be:

1. totally correlated with the discriminating variables y (x belongs to the set y for example). This is the case treated in Section 3.
2. uncorrelated with y . This is the subject of Section 4.
3. partly correlated with y . This case is discussed Section 8.

In an attempt to have access to the distributions of control variables, a common method consists of applying cuts which are designed to enhance the contributions to the data sample of particular sources of events. However, the result is frequently unsatisfactory: firstly because it can be used only if the signal has prominent features to be distinguished from the background, and secondly because of the cuts applied, a sizeable fraction of signal events can be lost, while a large fraction of background events may remain.

The aim of the *sPlot* formalism developed in this paper is to unfold the true distribution (denoted in boldface $\mathbf{M}_n(x)$) of a control variable x for events of the n^{th} species (any one of the N_s species), from the sole knowledge of the pdfs of the discriminating variables f_i , the first step being to proceed to the maximum Likelihood fit to extract the yields N_i . The statistical technique *sPlot* allows to build histograms in x keeping all signal events while getting rid of all background events, and keeping track of the statistical uncertainties per bin in x .

3 First step towards *sPlot*: *inPlot*

In this Section, as a means of introduction, one considers a variable x assumed to be totally correlated with y : x is a function of y . A fit having been performed to determine the yields N_i for all species, one can define naively, for all events, the weight

$$\mathcal{P}_n(y_e) = \frac{N_n f_n(y_e)}{\sum_{k=1}^{N_s} N_k f_k(y_e)}, \quad (2)$$

which can be used to build an estimate, denoted $\tilde{\mathbf{M}}_n$, of the x -distribution of the species labelled n (signal or background):

$$N_n \tilde{\mathbf{M}}_n(\bar{x}) \delta x \equiv \sum_{e \in \delta x} \mathcal{P}_n(y_e), \quad (3)$$

where the sum runs over the events for which the x value lies in the bin centered on \bar{x} and of total width δx .

In other words, $N_n \tilde{\mathbf{M}}_n(\bar{x}) \delta x$ is the x -distribution obtained by histogramming events, using the weight of Eq. (2). To obtain the expectation value of $\tilde{\mathbf{M}}_n$, one should replace the sum in Eq. (3) by the integral

$$\left\langle \sum_{e \in \delta x} \right\rangle \rightarrow \int dy \sum_{j=1}^{N_s} N_j f_j(y) \delta(x(y) - \bar{x}) \delta x. \quad (4)$$

Similarly, identifying the number of events N_i as determined by the fit to the expected number of events, one readily obtains:

$$\left\langle N_n \tilde{\mathbf{M}}_n(\bar{x}) \right\rangle \equiv N_n \mathbf{M}_n(\bar{x}). \quad (5)$$

Therefore, the sum over events of the naive weight \mathcal{P}_n reproduces, on average, the true distribution $\mathbf{M}_n(x)$. Plots obtained that way are referred to as *inPlots*: they provide a correct means to reconstruct $\mathbf{M}_n(x)$ only insofar as the variable considered is **in** the set of discriminating variables y . These *inPlots* suffer from a major drawback: x being fully correlated to y , the pdfs of x enter implicitly in the definition of the naive weight, and as a result, the $\tilde{\mathbf{M}}_n$ distributions cannot be used easily to assess the quality of the fit, because these distributions are biased in a way difficult to grasp, when the pdfs $f_i(y)$ are not accurate. For example, let us consider a situation where, in the data sample, some events from the n^{th} species show up far in the tail of the $\mathbf{M}_n(x)$ distribution which is implicitly used in the fit. The presence of such events implies that the true distribution $\mathbf{M}_n(x)$ must exhibit a tail which is not accounted for by $\tilde{\mathbf{M}}_n(x)$. These events would enter in the reconstructed *inPlot* $\tilde{\mathbf{M}}_n$ with a very small weight, and they would thus escape detection by the above procedure: $\tilde{\mathbf{M}}_n$ would be close to \mathbf{M}_n , the distribution assumed for x . Only a mismatch in the core of the x -distribution can be revealed with *inPlots*. Stated differently, the error bars which can be attached to each individual bin of

^bRemoving one of the discriminating variables from the set y before performing again the maximum Likelihood fit, one can consider the removed variable as a control variable x , provided it is uncorrelated with the others.

$\tilde{\mathbf{M}}_n$ cannot account for the systematical bias inherent to the *inPlots*.

4 The *sPlot* formalism

In this Section one considers the more interesting case where the two sets of variables x and y are uncorrelated. Hence, the total pdfs $f_i(x, y)$ all factorize into products $\mathbf{M}_i(x)f_i(y)$. While performing the fit, which relies only on y , no *a priori* knowledge of the x -distributions is used.

One may still consider the above distribution $\tilde{\mathbf{M}}_n$ (Eq. (3)), using the naive weight of Eq. (2). However in that case, the expectation value of $\tilde{\mathbf{M}}_n$ is a biased estimator of \mathbf{M}_n :

$$\begin{aligned} \langle N_n \tilde{\mathbf{M}}_n(\bar{x}) \rangle &= \int dy dx \sum_{j=1}^{N_s} N_j \mathbf{M}_j(x) f_j(y) \delta(x - \bar{x}) \mathcal{P}_n \\ &= N_n \sum_{j=1}^{N_s} \mathbf{M}_j(\bar{x}) N_j \int dy \frac{f_n(y) f_j(y)}{\sum_{k=1}^{N_s} N_k f_k(y)} \quad (6) \\ &\neq N_n \mathbf{M}_n(\bar{x}) . \end{aligned}$$

Here, the naive weight is no longer satisfactory because, when summing over the events, the x -pdfs $\mathbf{M}_j(x)$ appear now on the right hand side of Eq. (4), while they are absent in the weight. However, one observes that the correction term in the right hand side of Eq. (6) is related to the inverse of the covariance matrix, given by the second derivatives of $-\mathcal{L}$:

$$\mathbf{V}_{nj}^{-1} = \frac{\partial^2(-\mathcal{L})}{\partial N_n \partial N_j} = \sum_{e=1}^N \frac{f_n(y_e) f_j(y_e)}{(\sum_{k=1}^{N_s} N_k f_k(y_e))^2} . \quad (7)$$

On average, one gets:

$$\langle \mathbf{V}_{nj}^{-1} \rangle = \int dy \frac{f_n(y) f_j(y)}{\sum_{k=1}^{N_s} N_k f_k(y)} . \quad (8)$$

Therefore, Eq. (6) can be rewritten:

$$\langle \tilde{\mathbf{M}}_n(\bar{x}) \rangle = \sum_{j=1}^{N_s} \mathbf{M}_j(\bar{x}) N_j \langle \mathbf{V}_{nj}^{-1} \rangle . \quad (9)$$

Inverting this matrix equation, one recovers the distribution of interest:

$$N_n \mathbf{M}_n(\bar{x}) = \sum_{j=1}^{N_s} \langle \mathbf{V}_{nj} \rangle \langle \tilde{\mathbf{M}}_j(\bar{x}) \rangle . \quad (10)$$

Hence, when x is uncorrelated with the set y , the appropriate weight is not given by Eq. (2), but is

the covariance-weighted quantity (thereafter called *sWeight*) defined by:

$$\boxed{{}_s \mathcal{P}_n(y_e) = \frac{\sum_{j=1}^{N_s} \mathbf{V}_{nj} f_j(y_e)}{\sum_{k=1}^{N_s} N_k f_k(y_e)}} . \quad (11)$$

With this *sWeight*, the distribution of the control variable x can be obtained from the *sPlot* histogram:

$$N_n {}_s \tilde{\mathbf{M}}_n(\bar{x}) \delta x \equiv \sum_{e \in \delta x} {}_s \mathcal{P}_n(y_e) , \quad (12)$$

which reproduces, on average, the true binned distribution:

$$\langle N_n {}_s \tilde{\mathbf{M}}_n(x) \rangle = N_n \mathbf{M}_n(x) . \quad (13)$$

The fact that the covariance matrix \mathbf{V}_{ij} enters in the definition of the *sWeights* is enlightening: in particular, the *sWeight* can be positive or negative, and the estimators of the true pdfs are not constrained to be strictly positive.

5 *sPlot* properties

Beside satisfying the essential asymptotic property Eq. (13), *sPlots* bear properties which hold for finite statistics.

The distribution ${}_s \tilde{\mathbf{M}}_n$ defined by Eq. (12) is guaranteed to be normalized to unity and the sum over the species of the *sPlots* reproduces the data sample distribution of the control variable. These properties rely on maximizing the Likelihood:

- Each x -distribution is properly normalized. The sum over the x -bins of $N_n {}_s \tilde{\mathbf{M}}_n \delta x$ is equal to N_n :

$$\sum_{e=1}^N {}_s \mathcal{P}_n(y_e) = N_n . \quad (14)$$

- In each bin, the sum over all species of the expected numbers of events equals to the number of events actually observed. In effect, for any event:

$$\sum_{l=1}^{N_s} {}_s \mathcal{P}_l(y_e) = 1 . \quad (15)$$

Therefore, an *sPlot* provides a consistent representation of how all events from the various species are distributed in the control variable x . Summing up the N_s *sPlots*, one recovers the data sample distribution in x , and summing up the number of events entering in a *sPlot* for a given species, one recovers the yield of the species, as it is provided by the fit. For instance, if one observes an excess of events for a

particular n^{th} species, in a given x -bin, this excess is effectively accounted for in the number of events N_n resulting from the fit. To remove these events implies a corresponding decrease in N_n . It remains to gauge how significant is an anomaly in the x -distribution of the n^{th} species.

The statistical uncertainty on N_n ${}_s\tilde{M}_n(x)\delta x$ can be defined in each bin by

$$\sigma[N_n {}_s\tilde{M}_n(x)\delta x] = \sqrt{\sum_{e \in \delta x} ({}_s\mathcal{P}_n)^2}. \quad (16)$$

The above properties Eqs. (13)-(15) are completed by the fact that the sum in quadrature of the uncertainties Eq. (16) reproduces the statistical uncertainty on the yield N_n , as it is provided by the fit. In effect, the sum over the x -bins reads:

$$\sum_{[\delta x]} \sigma^2[N_n {}_s\tilde{M}_n\delta x] = \mathbf{V}_{nn}. \quad (17)$$

Therefore, for the expected number of events per x -bin indicated by the ${}_s\mathcal{P}$ lots, the statistical uncertainties are straightforward to compute using Eq. (16). The latter expression is asymptotically correct, and it provides a consistent representation of how the overall uncertainty on N_n is distributed in x among the events of the n^{th} species. Because of Eq. (17), and since the determination of the yields is optimal when obtained using a Likelihood fit, one can conclude that the ${}_s\mathcal{P}$ lot technique is itself an optimal method to reconstruct distributions of control variables.

6 Illustrations

An example of ${}_s\mathcal{P}$ lot at work is taken from the analysis where the method was first used ^{2,3}. One deals with a data sample in which three species are present: $B^0 \rightarrow \pi^+\pi^-$ and $B^0 \rightarrow K^+\pi^-$ are signals and the main background comes from $e^+e^- \rightarrow q\bar{q}$. The variable which is not incorporated in the fit is called ΔE and is used here as the control variable x . The detailed description of the variables can be found in Refs. ^{2,3}.

The left plot of Fig. 1 shows the distribution of ΔE after applying a cut on the Likelihood ratio. Therefore, the resulting data distribution concerns a reduced subsample for which statistical fluctuations cannot be attributed unambiguously to signal or to background. For example, the excess of events appearing on the left of the peak is likely to be attributed to a harmless background fluctuation.




Figure 1. Signal distribution of the ΔE variable. The left figure is obtained applying a cut on the Likelihood ratio to enrich the data sample in signal events (about 60% of signal is kept). The right figure shows the ${}_s\mathcal{P}$ lot for signal (all events are kept).

Looking at the right plot of Fig. 1, which is a signal ${}_s\mathcal{P}$ lot, one can see that these events are signal events, not background events. The pdf of ΔE which is used in the conventional fit for the whole analysis is superimposed on the ${}_s\mathcal{P}$ lot. When this pdf is used, the events in excess are interpreted as background events while performing the fit. Further studies have shown ² that these events are in fact radiative events, *i.e.* $B^0 \rightarrow \pi^+\pi^-\gamma$. When ignored in the analysis they lead to underestimates of the branching ratios by about 10%. The updated results ⁴ for the $B^0 \rightarrow \pi^+\pi^-$, $K^+\pi^-$ analysis, now taking into account the contribution of radiative events, show agreement with the estimate made in Ref. ².

7 Applications

Beside providing a convenient and optimal tool to cross-check the analysis by allowing distributions of control variables to be reconstructed and then compared with expectations, the ${}_s\mathcal{P}$ lot formalism can be applied also to extract physics results, which would otherwise be difficult to obtain. For example, one may be willing to explore some unknown physics involved in the distribution of a variable x . Or, one may be interested to correct a particular yield provided by the Likelihood fit from a selection efficiency which is known to depend on a variable x , for which the pdf is unknown. Provided one can demonstrate (*e.g.* through Monte-Carlo simulations) that the variable x exhibits weak correlation with the discriminating variables y .

To be specific, one can take the example of a three body decay analysis of a species, the signal,

polluted by background. The signal pdf inside the two-dimensional Dalitz plot is assumed to be not known, because of unknown contributions of resonances, continuum and of interference pattern. Since the x -dependence of the selection efficiency $\epsilon(x)$ can be computed without *a priori* knowledge of the x -distributions, one can build the efficiency corrected two-dimensional $sPlots$ (cf. Eq. (12)):

$$\frac{1}{\epsilon(\bar{x})} N_n \tilde{M}_n(\bar{x}) \delta x = \sum_{e \in \delta x} \frac{1}{\epsilon(x_e)} {}_sP_n(y_e), \quad (18)$$

and compute the efficiency corrected yields:

$$N_n^\epsilon = \sum_{e=1}^N \frac{{}_sP_n(y_e)}{\epsilon(x_e)}. \quad (19)$$

Analyses can then use the $sPlot$ formalism for validation purposes, but also, using Eq. (18) and Eq. (19), to probe for resonance structures and to measure branching ratios⁵.

8 Correlation between variables

Correlations between variables, if not trivial, are usually assessed by Monte-Carlo simulations. In case significant correlations are observed, one may still use the $sPlot$ weight of Eq. (11), but then there is a caveat. The distribution obtained with $sPlot$ cannot be compared directly with the marginal distribution of x . In that case, one must rely on Monte-Carlo simulation, and apply the $sPlot$ technique to the simulated events, in order to obtain Monte-Carlo $sPlots$. It is these Monte-Carlo $sPlots$ which are to be compared to the $sPlot$ obtained with the real data. Stated differently, the $sPlot$ can still be applied to compare the behaviour of the data with the Monte-Carlo expected behavior, but it loses its simplicity.

9 Conclusion

The technique presented in this paper applies when

- one examines a data sample originating from different sources of events,
- a Likelihood fit is performed on the data sample to determine the yields of the sources,
- this Likelihood uses a set y of discriminating variables,
- keeping aside a control variable x which is statistically uncorrelated to the set y .

By building $sPlots$, one can reconstruct the distributions of the control variable x , separately for each source present in the data sample. Although no cut is applied (hence, the $sPlot$ of a given species represents the whole statistics of this species) the distributions obtained are pure in a statistical sense: they are free from the potential background arising from the other species. The more discriminating the variables y , the clearer the $sPlot$ is. The technique is straightforward to implement; it is available in the ROOT framework under the class `TSPlot`⁶. It features several nice properties: both the normalizations and the statistical uncertainties of the $sPlots$ reflect the fit outputs.

References

1. M. Pivk and F.R. Le Diberder, *Nucl. Inst. Meth.A* 555, 356-369, 2005 (physics/0402083).
2. M. Pivk, Thèse de l'Université Paris VII, BABAR-THESIS-03/012 (2003), available (in French) at <http://tel.ccsd.cnrs.fr> (ID 00002991).
3. The BABAR Collaboration, *Phys. Rev. Lett.* **89**, 281802 (2002).
4. The BABAR Collaboration, hep-ex/0508046.
5. The BABAR Collaboration, *Phys. Rev. Lett.* **93**, 181805 (2004).
6. <http://root.cern.ch/root/html/doc/TSPlot.html>