G. Cowan, Unfolding blurb, v2.

1 Introduction

2 From Probability models to likelihoods

2.1 Basic concepts

2.1.1 Unfolded distributions

In some analyses the parameters of interest represent the expected numbers of entries in bins of a differential distribution. There are two basic approaches to this problem that we can call "unfolding" and "forward folding". The two methods lead to different requirements for what must be reported for further analysis.

Often when measuring a distribution one defines parameters $\vec{\mu} = (\mu_1, \ldots, \mu_M)$ to represent the expected number of entries in a bin assuming perfect resolution (so-called "particle-level" or "truth-level" parameters). In practice the real detector has limited resolution and so an event with a true value of a variable in a certain bin might be measured ("reconstructed") in a different one, so that $\vec{\nu} = (\nu_1, \ldots \nu_N)$ represents the expected numbers of events at reconstructed or "detector" level. These are related by

$$\vec{\nu} = R\vec{\mu}\,, \tag{1}$$

where R is an $N \times M$ *response matrix* defined such that R_{ij} represents the probability to be measured in bin i given that the true value was in bin j (here we neglect background processes).

Estimating the truth-level parameters $\vec{\mu}$ or *unfolding* of the distribution results in correlated estimators (see, e.g., Refs. [xxx]). The estimators are often treated as a Gaussian distributed vector characterized by an $M \times M$ covariance matrix $U_{ij} = \mathrm{cov}[\hat{\mu}_i, \hat{\mu}_j]$. In addition, the estimators are often constructed to include a small bias in exchange for a reduction in statistical variance (regularized unfolding).

By contrast, in forward folding one reports estimators for the expected numbers of events at detector level $\vec{\nu}$. In the simplest case one has $\hat{\nu}_i = n_i$, where n_i is the observed number of entries in the ith bin. Then to compare this result to the prediction of a certain model that predicts a particle-level distribution $\vec{\mu}$, one needs to "fold" the model prediction with the response matrix, i.e., one compares the $\vec{n} = (n_1, \ldots, n_N)$ to the $\vec{\nu}$ from Eq. (1). Often this likelihood will treat the n_i as independent and Poisson distributed.

The advantage of unfolding is that the estimated parameters represent directly the distribution in question. They can be compared between experiments and to model predictions using, e.g., a multivariate Gaussian likelihood (the covariance matrix of the estimators U is essential). Its disadvantage is that the unfolding may require regularization, which necessarily introduces some bias. In forward folding, one simply reports the observed numbers of events in the bins, and thus no regularization bias enters. But to compare these results with a model prediction, one needs the response matrix R.