HH

CERN-CN 93-3

828319

EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

9

CN/93/3

CM-P00065732

The Method of Reduced Cross-Entropy

A General Approach to Unfold Probability Distributions

Michael Schmelling / CERN

Abstract

A new unfolding method is presented for the case where the unfolded distribution is known to be non-negative everywhere. The method combines the least squares method with the principle of minimum cross-entropy. Its properties are discussed together with an algorithm for its realization and illustrated by means of a numerical example.

(Submitted to Nucl. Instrum. and Methods)

1 Introduction

Aside from purely statistical fluctuations, experimental measurements often are also systematically distorted with respect to the true distribution due to a non-uniform efficiency or the finite resolution of the apparatus. The deformation can be described by a convolution of the true distribution with the response function of the apparatus. Unfolding, i.e. correcting for those distortions, is comparatively easy if the true distribution is given by a parametric function with a small number of free parameters. Then a least squares fit can be used to infer the true distribution. This paper deals with the case that no such function is given. Typical examples include correcting momentum spectra, mass spectra or multiplicity distributions. An excellent introduction into the field can be found in reference [1].

Various approaches to solve the unfolding problem are possible. The most popular one relies on estimating a ratio between true and observed distribution and to use this ratio to correct the measurements. This approach has the distinct advantage of being simple and statistically stable, but is unsatisfactory because it requires an assumption about the true distribution which should be as close as possible to the truth [1]. It also is incapable to take into account information measured outside phase space limits.

Rigorous unfolding methods use only the known properties of the apparatus in order to obtain an estimate for the true distribution. Using for instance histograms to represent the true and the observed distribution, the convolution which describes the measurements can be approximated by a matrix equation. The observed distribution is obtained as the product of the true distribution with the response matrix of the apparatus, which may be known from test measurements or detailed Monte-Carlo simulations. Of course one could try to infer the true distribution by applying the inverse of the response matrix to the measurements. However, due to the statistical errors of the measurements this tends to yield extremely unstable results [1]. Because of these instabilities the simple method described before usually is the preferred method, despite its conceptual weakness.

An alternative is offered by regularization methods, which combine the rigorous approach with an attempt to filter out the statistical fluctuations. Then the true distribution for example is estimated by the smoothest or least curved distribution which is still statistically compatible with the measurements [1].

The “Method of Reduced Cross-Entropy” (MRX) described here is a regularization method which is applicable whenever the true distribution is known to be non-negative everywhere. It is based on the result [2, 3], that the unique consistent estimate for a probability distribution which is not specified unambiguously by a set of measurements is given by the principle of minimum cross-entropy [4]. An introduction to the principle of minimum cross-entropy is given in appendix A.

After formally stating the general linear unfolding problem in section 2, the MRX and its properties are discussed in section 3. Section 4 illustrates its practical application by means of a numerical example.

2 The General Linear Unfolding Problem

The MRX as described in this paper will be formulated for the discrete case of the general linear unfolding problem, where the true distribution is the probability distribution $\bar{p}_j$

$$\bar{p}_j > 0, \quad j = 1 \dots n \quad \text{and} \quad \sum_{j=1}^n \bar{p}_j = 1. \quad (1)$$

A set of measurements a_i^{obs} providing experimental information about $\bar{p}_j$ then can be written as

$$a_i^{obs} = \sum_{j=1}^n G_{ij} \bar{p}_j + \Delta_i = \bar{a}_i + \Delta_i \quad \text{with} \quad i = 1 \dots m. \quad (2)$$

In this expression the response matrix G_{ij} describes how the apparatus distorts the true distribution. Ideally one would have $G_{ij} = \delta_{ij}$. The offsets Δ_i are statistical fluctuations by which the actual measurements a_i^{obs} deviate from their expectation values $\bar{a}_i$.

For the following it will be assumed that the response matrix G_{ij} is known together with the measurements a_i^{obs} and their covariance matrix C_{ik}

$$C_{ik} = \langle (a_i^{obs} - \bar{a}_i)(a_k^{obs} - \bar{a}_k) \rangle = \langle \Delta_i \Delta_k \rangle. \quad (3)$$

The general linear unfolding problem then can be stated as follows:

Given is a set of measurements $a_i^{obs}, i = 1 \dots m$ with covariance matrix C_{ik} and the response matrix G_{ij} which relates those measurements to an unknown probability distribution $\bar{p}_j, j = 1 \dots n$. Based only on this information, the aim is to derive an estimate p_j of the unknown distribution $\bar{p}_j$ for arbitrary values m and n .

This formulation is general in the sense that it admits arbitrary dimensions m and n . Its solution requires a method of inference which is able to provide a meaningful estimate of the true distribution both when $\bar{p}_j$ is overconstrained ($m > n$) and when it is underconstrained ($m < n$) by the measurements. Optimal solutions for limiting cases of both classes of unfolding problems can be found in the literature. For the case of the overconstrained problem it is the least squares method [5], which yields an unbiased estimate with minimum variance. For the underconstrained problem and disregarding statistical errors the unique consistent estimate for a probability distribution is obtained by the principle of minimum cross-entropy [2, 3].

3 The Method of Reduced Cross-Entropy

The MRX presented in this paper is a synthesis of the principle of minimum cross-entropy with the least squares method. It extends the principle of minimum cross-entropy to the case that the measurements are subject to statistical fluctuations and the least squares method to deal with underconstrained problems. The current section focuses on a comprehensive presentation of the MRX and its properties. Technical details and mathematical proofs are given in the appendices.

3.1 Definition of the Method of Reduced Cross-Entropy

The MRX determines the solution of the general linear unfolding problem as that probability distribution p_j which minimizes

$$F = w\chi^2 + S, \quad w = \text{const} > 0, \quad (4)$$

where

$$\chi^2 = \sum_{i,k=1}^m (a_i - a_i^{obs}) C_{ik}^{-1} (a_k - a_k^{obs}) \quad \text{with} \quad a_i = \sum_{j=1}^n G_{ij} p_j \quad \text{and} \quad S = \sum_j p_j \ln \frac{p_j}{\varepsilon_j}. \quad (5)$$

Here χ^2 measures the discrepancy between the estimate p_j and the measurements a_i^{obs} and the cross-entropy S the deviation of this estimate from the “prior distribution” ε_j which contains any a priori knowledge about the true distribution. If nothing is known except the fact that the true distribution is positive everywhere, one has $\varepsilon_j = \text{const}$. In this case the MRX becomes the “Method of Reduced Entropy” [6]. The “regularization parameter” w balances a better agreement with the measurements (smaller χ^2) against a closer match with the prior distribution (smaller S). How to determine its numerical value will be discussed later. For a given regularization parameter w the minimum of (4) specifies the distribution p_j uniquely. The proof is given in appendix (B) together with an algorithm to actually determine p_j .

Following the arguments presented in appendix (A) equation (4) can be interpreted as a maximum likelihood ansatz, where the logarithm of the statistical likelihood (χ^2) for an estimate p_j is combined with the logarithm of the likelihood that it is compatible with the prior distribution ε_j .

How the ansatz (4) unifies the least squares method with the principle of minimum cross-entropy becomes most transparent in the limit $w \rightarrow \infty$, or equivalently when the statistical errors of the measurements go to zero. In this case, since S is bounded both from above and from below, the function F will always be dominated by the χ^2 -term. If the problem is overconstrained the estimate p_j will be determined uniquely by the solution of minimum χ^2 , i.e. according to the least squares method. If the problem is underconstrained the requirement of minimum χ^2 does not specify a unique solution p_j . The minimum of F then will be obtained by that distribution p_j which for minimum χ^2 has the smallest cross-entropy S .

3.2 The Regularization Matrix

The interplay between χ^2 and S in the minimization of F can be seen most clearly when starting at the minimum value of χ^2 . Here $w\chi^2$ varies quadratically while S in general will vary linearly. For the minimization of F a slight increase in $w\chi^2$ thus will be more than compensated by a corresponding decrease of S . Any deviation from the minimum χ^2 can be viewed as ignoring part of the experimental information. How much information is retained can be formulated quantitatively by means of the “regularization matrix” R_{ij}

$$R_{ij} = \frac{\partial a_j}{\partial a_i^{obs}} \quad (6)$$

which describes how the estimates a_i (eq.5) are coupled to the measurements a_i^{obs} . The trace of this matrix can be interpreted as the effective number of measurements actually used to determine the estimate p_j

$$M_{eff} = \sum_{i=1}^m R_{ii}. \quad (7)$$

Functional form and specific properties of the regularization matrix are derived in appendix C. For finite values w the quantity M_{eff} is a continuous variable with $0 \leq M_{eff} \leq m$. In the limit $w \rightarrow \infty$ it converges towards an integer. Generalizing the standard χ^2 -test M_{eff} can be employed to determine a χ^2 -confidence level for the unfolding result. With the number of degrees of freedom given by $N_{df} = m - M_{eff}$ this confidence level is obtained in the usual way from the incomplete gamma function Q

$$CL = Q(N_{df}/2, \chi^2/2) = \frac{1}{\Gamma(\frac{N_{df}}{2})} \int_{\frac{\chi^2}{2}}^{\infty} dt e^{-t} t^{\frac{N_{df}}{2}-1} \quad (8)$$

which is well defined for both continuous χ^2 and N_{df} . Small confidence levels either indicate that the regularization parameter w was chosen too small or that the response matrix fails to describe the apparatus.

3.3 The Choice of the Regularization Parameter

Deviating from the least squares solution, the unfolding result is always slightly biased towards the prior distribution. The size of this bias is controlled by the regularization parameter w . A variety of criteria which yield reasonable values for the regularization parameter is conceivable and there is a certain freedom which one to choose as long as the unfolding result describes the measurements with an acceptable confidence level (8). Then the bias can be expected to be small compared to the statistical errors. The criterion put forward here requires w to be chosen such that an estimated deviation of the actual χ^2 from its minimum value is one unit:

$$\Delta\chi^2 = \sum_{ijkl=1}^m (a_k - a_k^{obs}) R_{ki} C_{ij}^{-1} (a_l - a_l^{obs}) R_{lj} = 1. \quad (9)$$

The expression for $\Delta\chi^2$ is derived from the χ^2 definition 5 by weighting the residuals with the coupling of the a_i to the measurements. Those components of a_i which decouple from the measurements and thus in the limit $w \rightarrow \infty$ account for the asymptotic χ^2 do not contribute. The weighting with R_{ik} is a way to suppress the χ^2 -offset. Further details are discussed in appendix D.

3.4 Invariance Under Regular Transformations

One important aspect of any consistent unfolding method is the invariance of the result with respect to any regular linear transformation of the constraints. Since the information content of the measurements is not changed by such a transformation, the unfolding result must be invariant too. The MRX satisfies this requirement. As shown in appendix E the unfolded distribution p_j is invariant. The same holds for the regularization parameter w , the number M_{eff} of measurements effectively used in the unfolding and the χ^2 of the final result, which gives another justification for the goodness-of-fit test according to equation (8).

4 A Numerical Example

Using the example given already in reference [1] the MRX is employed to correct the measurement of a one dimensional distribution for distortions due to detector effects¹. The true distribution is taken to be

$$\bar{p}(y) = \frac{4}{4 + (y - 0.4)^2} + \frac{0.4}{0.04 + (y - 0.8)^2} + \frac{0.2}{0.04 + (y - 1.5)^2} \quad \text{with} \quad 0 \leq y \leq 2, \quad (10)$$

and the response function which maps the true to the observable distribution as

$$g(x, y) = \left[1 - \frac{1}{2}(1 - y)^2 \right] \left[\frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x - y + 0.05 y^2)^2}{2\sigma^2}\right) \right] \quad \text{with} \quad \sigma = 0.1. \quad (11)$$

The first term in square brackets describes a detection efficiency which varies between 0.5 and 1, the second one a resolution function which is taken to be a simple gaussian smearing with the observed variable x biased proportional to y^2 . The relation between true and observed distribution is

$$\bar{a}(x) = \int_0^2 dy g(x, y) \bar{p}(y). \quad (12)$$

In order to apply the algorithms developed in this paper a discrete approximation of equation (12) needs to be used. To go from the probability density $\bar{p}(y)$ to a probability distribution $\bar{p}_j$, the true distribution is integrated over finite y -intervals. Applying the same procedure to $\bar{a}(x)$ yields

$$\bar{a}_i = G_{ij} \bar{p}_j \quad (13)$$

with

$$\bar{a}_i = \int_{x_{i-1}}^{x_i} dx \bar{a}(x), \quad \bar{p}_j = \int_{y_{j-1}}^{y_j} dy \bar{p}(y) \quad (14)$$

and

$$G_{ij} = \frac{\int_{x_{i-1}}^{x_i} dx \int_{y_{j-1}}^{y_j} dy g(x, y) \bar{p}(y)}{\int_{y_{j-1}}^{y_j} dy \bar{p}(y)} \approx \frac{1}{y_j - y_{j-1}} \int_{x_{i-1}}^{x_i} dx \int_{y_{j-1}}^{y_j} dy g(x, y). \quad (15)$$

From (15) one sees that a reasonable approximation of the continuous problem is obtained as soon as the bin width for the true distribution becomes smaller than the width of the resolution function. Then the response matrix G_{ij} becomes practically independent of the true distribution. Choosing sufficiently narrow bins the original problem can be approximated to any desired precision.

For the following the observed distribution $\bar{a}(x)$ is considered over the range $0 \leq x \leq 2$. Equal size bins are used to represent both distributions, $n = 20$ for the true and three values $m = 10, 20$ and 40 for the observed one. Figure 1a displays the true distribution compared to a histogram of the distribution $\bar{a}_i$. Simulating actual measurements the $\bar{a}_i$ are fluctuated with the size of the fluctuations corresponding to the result of a counting experiment which recorded 10^4 events. These “measurements” a_i^{obs} are represented by the points with error bars.

Unfolding the measurements now proceeds in two steps. In the first step the MRX is used to unsmear the normalized measurements, and in the second step the result of the unsmearing is

¹A FORTRAN implementation of the MRX together with this example is available from the author.

properly re-normalized and efficiency corrected. The choice for the prior distribution ε_j to be used in the unsmearing step is chosen according to the principle of minimum cross-entropy. Given no additional input for the true distribution is given by the “phase space” prediction, i.e. if nothing else is known it is taken to be flat. Then the prior distribution for the unsmearing step becomes $\varepsilon_j \sim (y_j - y_{j-1}) \cdot \eta_j$, where η_j denotes the detection efficiency in bin j of the true distribution.

The unfolding result obtained with the MRX compared to the true distribution is shown in Fig.1b. Also shown in the inset is the result one obtains by simply inverting the response matrix. The oscillating pattern is a consequence of the strong correlations between adjacent bins. In addition there are huge error bars. The dotted line shows the true distribution whose amplitude of variation turns out to be negligible compared to the size of the errors of the corrected distribution. Though statistically correct, this result apparently holds no information about the shape of the true distribution. It even violates the a priori knowledge that the true distribution is non-negative everywhere. Only for much smaller measurement errors than assumed in this example the result from the simple matrix inversion stabilizes. The comparison between the MRX result and what one gets from an inversion of the response matrix thus underlines the importance of filtering statistical fluctuations in order to arrive at meaningful results.

The concept of regularization versus matrix inversion can also be interpreted in the following way [7]: The inversion method essentially is an attempt to undo all detector effects up to the resolution implied by the bin size chosen for the true distribution. The regularization approach corrects up to a residual smearing which is determined by the size of the statistical errors of the measurements. The width of the residual smearing can be extracted quantitatively from the covariance matrix of the unfolded distribution which contains the complete information about how the bins are correlated. Since the rms-width of a rectangular distribution is $1/\sqrt{12}$ times its total width, the rms-width of the residual smearing in units of bins can be estimated by

$$\sigma_{res} = \frac{n}{M_{eff}} \frac{1}{\sqrt{12}}. \quad (16)$$

In the current example with $n = 20$ bins for the unfolded distribution one finds $M_{eff} = 11.2$, which corresponds to a residual resolution of $\sigma_{res} = 0.5$ bins. This is only half the size of the original smearing. The unfolded result thus not only is bias and efficiency corrected, but is also able to resolve structures which are significantly narrower than those which are visible in the uncorrected measurements. Comparing the “observed” distribution from Fig.1a to the unfolding result in Fig.1b one sees that the price to pay for this improvement is that the error bars of the unfolded distribution are larger than those of the measurements. In other words, 10^4 events observed with a given resolution function correspond to fewer events at an improved resolution.

Figure 2 shows how in the limit of small statistical errors the MRX converges towards the least squares method or the principle of minimum cross-entropy. Using 20 bins to represent the true distribution, Fig.2a shows the MRX estimate where 10 bins were used to represent the measured distribution compared to the result obtained according to the principle of minimum cross-entropy. Both agree perfectly. A very close inspection reveals that the peaks in the unfolded distribution are slightly broader than in the true distribution. No such bias of the unfolding results relative to the true distribution is visible in Fig.2b where 40 bins were used to represent the measurements. Here the MRX result is compared to the least squares estimate and again found to be in very close agreement.

5 Summary

The MRX has been shown to be a universal method of inference for a probability distribution which is constrained by a set of measurements. It is a synthesis of two methods which are known to be optimal in certain limiting cases: the least squares method and the principle of minimum cross-entropy. The MRX complements the principle of minimum cross-entropy with the ability to handle statistical fluctuations and the least squares method with a way of dealing with underconstrained problems. It has a unique solution and is invariant with respect to any regular linear transformations of the measurements. The compatibility of the unfolding result with the measurements can be judged by means of a generalized χ^2 -test.

Acknowledgements

It is a pleasure to thank Rafel Carreras for unknowingly pointing out to me that the method of maximum entropy can be employed to solve unfolding problems, Daniel Drijard for many useful discussions which lead to the insight that both χ^2 and entropy are related to probability concepts, ALEPH and CERN for its stimulating atmosphere and certainly Dietrich Wegener for his constant support during all steps in the evolution of the method of reduced cross-entropy.

Appendix

The following sections focus on specific technical and mathematical aspects concerning the method of reduced cross-entropy. Even though some of the points discussed here are textbook material, they are kept in order to provide a coherent presentation of the method. Depending on what is more convenient, equations are either written in matrix or index notation. Where appropriate Einstein's summation convention is adopted. Normal characters are used to denote scalar quantities or individual matrix elements, bold face characters refer to entire arrays or matrices. A bold-face **1** denotes the unit matrix.

A The Principle of Minimum Cross-Entropy

The principle of minimum cross-entropy specifies a method how to infer an estimate p_j for an unknown probability distribution $\bar{p}_j$ when the available information is insufficient to specify it uniquely. If this information is given in terms of a set of expectation values

$$\bar{a}_i = \sum_{j=1}^n G_{ij} \bar{p}_j, \quad \text{with} \quad i = 1 \dots m < n \quad (17)$$

this estimate is that distribution p_j out of all solutions satisfying the constraints (17) which minimizes the cross-entropy S

$$S = \sum_j p_j \ln \frac{p_j}{\varepsilon_j}. \quad (18)$$

Generalizing Shannon's "Information Entropy" [8], S is a measure for the difference between the distribution p_j and a "prior distribution" ε_j which contains any a priori knowledge about the true distribution. If nothing is known except the fact that the true distribution is non-negative everywhere one has $\varepsilon_j = \text{const}$, i.e. a priori all probabilities of the true distribution are taken to be equal. The principle of minimum cross-entropy then becomes identical to the principle of maximum entropy [9].

A rigorous axiomatic derivation of the principle of minimum cross-entropy can be found in [2]. Here a more intuitive explanation shall be given, which shows that the principle of minimum cross-entropy can be viewed as a maximum likelihood estimate. Out of all possible solutions p_j satisfying the constraints (17) that one is selected, which has the largest likelihood of being compatible with the prior distribution ε_j . The following considerations show that up to a constant term the logarithm of this likelihood is proportional to $-S$.

Having N events distributed according to the estimated probability distribution p_j one expects $n_j = N p_j$ events in each class. The likelihood $\mathcal{L}$ that a distribution of n_j events is observed if the true probabilities are given by ε_j is given by the multinomial distribution

$$\mathcal{L} = \frac{N!}{n_1! n_2! \dots n_n!} \varepsilon_1^{n_1} \varepsilon_2^{n_2} \dots \varepsilon_n^{n_n}. \quad (19)$$

Taking the logarithm of $\mathcal{L}$ and approximating the factorials by the leading term of Stirling's formula

$$\ln x! \approx x(\ln x - 1) + \frac{1}{2} \ln(2\pi x) \approx x(\ln x - 1) \quad (20)$$

one gets

$$\ln \mathcal{L} = \ln N! + \sum_{i=1}^n \ln \frac{\varepsilon_i^{n_i}}{n_i!} \approx \ln N! + \sum_{i=1}^n [n_i \ln \varepsilon_i - n_i (\ln n_i - 1)] = \text{const} - \sum_{i=1}^n n_i \ln \frac{n_i}{\varepsilon_i}, \quad (21)$$

and after re-substituting $n_i = Np_i$ one obtains

$$\ln \mathcal{L} \approx \text{const} - N \sum_{i=1}^n p_i \ln \frac{p_i}{\varepsilon_i} = \text{const} - NS. \quad (22)$$

B Finding the Distribution of Reduced Cross-Entropy

The problem of finding the distribution p_j of reduced cross-entropy which minimizes F (equation (4)) can be solved by generalizing the approach [10] for finding the distribution of maximum entropy. The solution is obtained by minimizing the function P

$$P = \frac{1}{4w} \lambda_k \lambda_l C_{kl} - \lambda_k a_k^{obs} + \ln \sum_{i=1}^n \beta_i \quad \text{where} \quad \beta_i = \varepsilon_i e^{\lambda_k G_{ki}} \quad (23)$$

with respect to the parameters λ_i . For each set of parameters λ_i a probability distribution is defined by

$$p_i = \beta_i / \sum_j \beta_j. \quad (24)$$

which becomes the distribution of reduced cross-entropy at the minimum of P . The proof proceeds in two steps, the first one establishing that P has a unique minimum and the second one showing that this minimum also determines the minimum of F .

Since C_{kl} is positive definite the quadratic term $\lambda_k \lambda_l$ which dominates the behavior of P for $\lambda_k \rightarrow \pm\infty$ is bounded from below. All other terms do not vary faster than linearly with λ_i . Consequently P as a whole is bounded from below, i.e. it must have a global minimum. The derivatives of P with respect to the λ_i are given by

$$P'_k = \frac{1}{2w} C_{kl} \lambda_l - a_k^{obs} + a_k \quad \text{with} \quad a_k = \sum_{j=1}^n G_{kj} p_j, \quad \text{and} \quad (25)$$

$$P''_{kl} = \frac{1}{2w} C_{kl} + \frac{\partial a_k}{\partial \lambda_l} = \frac{1}{2w} C_{kl} + Q_{kl} \quad \text{with} \quad Q_{kl} = \sum_i G_{ki} G_{li} p_i - \sum_i G_{ki} p_i \sum_j G_{lj} p_j. \quad (26)$$

Since $p_j > 0$ for any set of parameters λ_i the matrix Q_{kl} is always positive definite. By definition the same holds for C_{kl} . Thus also the second derivative of P is always positive definite, i.e. the function P is concave everywhere. It follows that the global minimum is the unique minimum.

With P equation (4) can be rewritten as

$$F = w\chi^2 - S = wP'_k P'_l C_{kl}^{-1} - P \quad (27)$$

and the condition for F to acquire a minimum becomes

$$F'_m = 2wP'_l (Q_{mk} C_{kl}^{-1}) = 0. \quad (28)$$

One solution is given by the minimum of P where P'_m vanishes. That this is the only solution where the gradient of F vanishes can be seen from the fact that, due to the positive definiteness of Q_{km} and C_{kl} each non-zero gradient P'_l results in the norm of F'_m being different from zero. Numerically the minimum of P can be found by a simple Gauss-Newton minimization procedure.

The covariance matrix $C(p)_{ij}$ of the distribution p_j is given by

$$C(p)_{ij} = \frac{\partial p_i}{\partial a_k^{obs}} \frac{\partial p_j}{\partial a_l^{obs}} C_{kl} \quad (29)$$

where, using the result (34) derived below, the partial derivatives are

$$\frac{\partial \ln p_i}{\partial a_k^{obs}} = \frac{\partial}{\partial a_k^{obs}} \left[\ln(\beta_i) - \ln\left(\sum_l \beta_l\right) \right] = \frac{\partial \lambda_j}{\partial a_k^{obs}} G_{ji} - p_l G_{jl} \frac{\partial \lambda_j}{\partial a_k^{obs}} = (P'')^{-1}_{jk} (G_{ji} - a_j). \quad (30)$$

C Properties of the Regularization Matrix

To determine the functional form of the regularization matrix, first the sensitivity of the λ_k to changes of the measurements a_i^{obs} shall be determined. Starting from the condition $P'_i = 0$ one has

$$2w(a_k - a_k^{obs})C_{ki}^{-1} + \lambda_i = 0. \quad (31)$$

Considering infinitesimal changes in the measurements yields

$$2w(da_k - da_k^{obs})C_{ki}^{-1} + d\lambda_i = 0, \quad (32)$$

which, using $da_k = (\partial a_k / \partial \lambda_l) d\lambda_l$, can be rewritten as

$$da_m^{obs} = \left(\frac{\partial a_m}{\partial \lambda_l} + \frac{1}{2w} C_{ml} \right) d\lambda_l. \quad (33)$$

The terms in parenthesis can be identified with the second derivative of P and one obtains

$$\frac{\partial \lambda_k}{\partial a_i^{obs}} = (P'')^{-1}_{ki}. \quad (34)$$

From this the regularization matrix can be read off as

$$R_{ij} = \frac{\partial a_j}{\partial a_i^{obs}} = \frac{\partial a_j}{\partial \lambda_k} \frac{\partial \lambda_k}{\partial a_i^{obs}} = \left(P''_{jk} - \frac{1}{2w} C_{jk} \right) (P'')^{-1}_{ki} = \delta_{ij} - \frac{1}{2w} C_{jk} (P'')^{-1}_{ki}. \quad (35)$$

As shown in appendix (E) the MRX is invariant under arbitrary regular transformations of the measurements. Exploiting this invariance, the properties of the regularization matrix are most conveniently discussed for a basis where the covariance matrix becomes the unit matrix and, using the additional freedom of another orthogonal transformation, where $\mathbf{Q}$ as introduced in (26) is diagonal. Then also the regularization matrix is diagonal with

$$R_{ii} = \frac{2wQ_{ii}}{1 + 2wQ_{ii}}, \quad i = 1 \dots m. \quad (36)$$

Since the eigenvalues of $\mathbf{Q}$ all are non negative, it follows that the diagonal elements of $\mathbf{R}$ are numbers between zero and one. For small values of w the regularization matrix is proportional to w . For $w \rightarrow \infty$ the behavior of $\mathbf{R}$ depends on the rank of $\mathbf{Q}$, which is given by

$$\text{Rank}(\mathbf{Q}) = \min(m, n_{pos}) \quad (37)$$

with m the number of measurements and n_{pos} the number of probabilities in $p_j, j = 1 \dots n$ which are greater than zero. At finite values w one has $\text{Rank}(\mathbf{Q}) = \min(m, n)$. In the limit $w \rightarrow \infty$ some of the probabilities p_j may go to zero, resulting in a reduced rank for $\mathbf{Q}$. In this limit the matrix elements (36) are either zero or one, with the number of ones equal to $\text{Rank}(\mathbf{Q})$.

D Properties of the Regularization Parameter

This section complements subsection (3.3) by giving an alternative derivation of the criterion (9) which defines the regularization parameter together with a more detailed discussion of the properties of this criterion.

As a first step a differential equation relating the parameter vector λ_k to changes of w shall be derived. Starting again from the condition $P'_i = 0$ one has

$$\lambda_m C_{ml} = 2w(a_l^{obs} - a_l) \quad (38)$$

and considering infinitesimal changes dw yields

$$d\lambda_m C_{ml} = 2dw(a_l^{obs} - a_l) - 2w da_l = \frac{dw}{w} \lambda_m C_{ml} - 2w da_l \quad (39)$$

which, using $da_l = (\partial a_l / \partial \lambda_m) d\lambda_m$, can be rewritten as

$$d\lambda_m \left(\frac{1}{2w} C_{ml} + \frac{\partial a_l}{\partial \lambda_m} \right) = \frac{dw}{w} \lambda_m \frac{1}{2w} C_{ml}. \quad (40)$$

Identifying the terms in parenthesis with the second derivative of P one obtains

$$\frac{\partial \lambda_k}{\partial w} = \frac{1}{2w^2} \lambda_m C_{ml} (P'')_{lk}^{-1} = \frac{1}{w} (\delta_{km} - R_{km}) \lambda_m. \quad (41)$$

Employing the properties of the regularization matrix $\mathbf{R}$ derived in appendix C the asymptotic behavior of λ_k can be determined from (41). For $w \rightarrow 0$ one finds

$$\frac{\partial \lambda_k}{\partial w} = \frac{\lambda_k}{w} \quad \Rightarrow \quad \lambda_k \sim w. \quad (42)$$

The behavior in the limit $w \rightarrow \infty$ again is most conveniently discussed for the diagonal form of the regularization matrix introduced in appendix C. For those parameters λ_k with $R_{km} \rightarrow \delta_{km}$ one gets

$$\frac{\partial \lambda_k}{\partial w} \rightarrow 0 \quad \Rightarrow \quad \lambda_k \rightarrow \text{const.} \quad (43)$$

The ones with $R_{km} \rightarrow 0$ diverge proportionally to w

$$\frac{\partial \lambda_k}{\partial w} \rightarrow \frac{\lambda_k}{w} \quad \Rightarrow \quad \lambda_k \sim w. \quad (44)$$

With the information gathered so far it is now possible to construct an estimate for the χ^2 -bias introduced by a given choice of w . Using equation (38), one obtains for χ^2

$$\chi^2 = (a_k^{obs} - a_k)(a_l^{obs} - a_l)C_{kl}^{-1} = \frac{1}{4w^2} \lambda_k \lambda_l C_{kl}. \quad (45)$$

One sees, that in the limit $w \rightarrow \infty$ the minimal χ^2 is due to those parameters λ_k which diverge proportionally to w . An estimate for the χ^2 -bias is obtained by filtering out the divergent component of the λ_k according to

$$\Delta \chi^2 = \frac{1}{4w^2} \lambda_k \left(1 - \frac{\partial \ln |\lambda_k|}{\partial \ln w}\right) \lambda_l \left(1 - \frac{\partial \ln |\lambda_l|}{\partial \ln w}\right) C_{kl} = \frac{1}{4w^2} \left(\lambda_k - w \frac{\partial \lambda_k}{\partial w}\right) \left(\lambda_l - w \frac{\partial \lambda_l}{\partial w}\right) C_{kl} \quad (46)$$

which, using (41), reduces to

$$\Delta \chi^2 = \frac{1}{4w^2} R_{km} \lambda_m R_{ln} \lambda_n C_{kl}. \quad (47)$$

From equation (35) follows that the product $\mathbf{RC}$ is a symmetric matrix. Using this information together with equation (38) it is easy to show that (9) and (47) are identical. With

$$\frac{1}{2w} C_{ik} R_{km} \lambda_m = \frac{1}{2w} \lambda_m C_{mk} R_{ki} = (a_k^{obs} - a_k) R_{ki} \quad (48)$$

one immediately finds

$$\Delta \chi^2 = \frac{1}{4w^2} R_{km} \lambda_m R_{ln} \lambda_n C_{kl} = \frac{1}{4w^2} C_{ik} R_{km} \lambda_m C_{jl} R_{ln} \lambda_n C_{ij}^{-1} = (a_k^{obs} - a_k) R_{ki} (a_l^{obs} - a_l) R_{lj} C_{ij}^{-1} \quad (49)$$

Finally the w -dependence of $\Delta \chi^2$ needs to be discussed. This is being done most conveniently using equation (47). In the limit $w \rightarrow 0$ both the parameters λ_k and the regularization matrix R_{ik} are proportional to w . Introducing σ to denote the scale of the measurement errors one obtains $\Delta \chi^2 \sim w^2 \sigma^2$. For $w \rightarrow \infty$ the diverging components of the parameters λ_k are filtered out by the regularization matrices. Only the constant components remain. Thus asymptotically one has $\Delta \chi^2 \sim w^{-2} \sigma^2$. As a consequence there will be two solutions w which satisfy the requirement $\Delta \chi^2 = 1$, one on the rising edge $\sim w^2$ and one at the falling edge $\sim w^{-2}$. Which one to take can be decided from considering how the solution scales with the measurement errors. On the rising edge the requirement $\Delta \chi^2 = 1$ results in $w \sim 1/\sigma$, on the falling edge one gets $w \sim \sigma$. According to equation (38) the difference between the measurements and the estimates obtained by the MRX scale like

$$a_i^{obs} - a_i \sim \sigma^2 / w. \quad (50)$$

Since those differences should scale with the measurement errors, one needs $w \sim \sigma$, i.e. the solution on the falling edge of $\Delta \chi^2$ has to be taken.

Figure 3 illustrates how, for the unfolding example given in Fig.1, the estimate $\Delta \chi^2$ varies as function of the regularization parameter w . The uncertainty due to statistical fluctuations is indicated by the width of the band which is small for most values of w . One sees that the criterion $\Delta \chi^2 = 1$ selects a value w just at the edge where $\Delta \chi^2$ becomes unstable, thereby ensuring that no significant information is lost while still maintaining an efficient filtering of statistical fluctuations.

E Transformation Properties of the MRX

Let $\mathbf{M}$ be any regular $[m \times m]$ matrix. Denoting transformed quantities by a tilde the unfolding problem can be re-expressed in terms of

$$\tilde{\mathbf{a}}^{obs} = \mathbf{M}\mathbf{a}^{obs}, \quad \tilde{\mathbf{C}} = \mathbf{M}\mathbf{C}\mathbf{M}^T \quad \text{and} \quad \tilde{\mathbf{G}} = \mathbf{M}\mathbf{G}. \quad (51)$$

The invariance of χ^2 follows immediately from these relations.

$$\tilde{\chi}^2 = (\tilde{\mathbf{a}}^{obs})^T (\tilde{\mathbf{C}})^{-1} (\tilde{\mathbf{a}}^{obs}) = (\mathbf{a}^{obs})^T \mathbf{M}^T (\mathbf{M}^T)^{-1} \mathbf{C}^{-1} \mathbf{M}^{-1} \mathbf{M} \mathbf{a}^{obs} = (\mathbf{a}^{obs})^T \mathbf{C}^{-1} \mathbf{a}^{obs} = \chi^2. \quad (52)$$

The function P , the minimum of which determines the solution of the unfolding problem, becomes

$$P = \frac{1}{4w} \lambda_k \lambda_l \tilde{C}_{kl} - \lambda_k \tilde{a}_k^{obs} + \ln \sum_{i=1}^n \varepsilon_i e^{\lambda_k \tilde{G}_{ki}} = \frac{1}{4w} \lambda_k \lambda_l M_{lj} C_{ji} M_{ik} - \lambda_k M_{kl} a_l^{obs} + \ln \sum_{i=1}^n \varepsilon_i e^{\lambda_k M_{ki} G_{li}}. \quad (53)$$

Thus, if the parameters vector λ_k determined the solution of the original unfolding problem, the transformed problem is solved by $M_{kl}^{-1} \lambda_l$ which leaves the distribution of reduced cross-entropy p_j unchanged.

The transformation law for the regularization matrix follows from its definition (6). One finds

$$\tilde{R}_{ml} = \frac{d\tilde{a}_l}{d\tilde{a}_m^{obs}} = \frac{da_k}{da_j^{obs}} M_{lk} M_{jm}^{-1} = M_{jm}^{-1} R_{jk} M_{lk} \quad \text{or} \quad \tilde{\mathbf{R}} = (\mathbf{M}^{-1})^T \mathbf{R} \mathbf{M}^T. \quad (54)$$

As a consequence M_{eff} , the effective number of measurements used to infer the unfolded distribution and thereby also the number of degrees of freedom $N_{df} = m - M_{eff}$ is invariant:

$$\tilde{M}_{eff} = \text{Tr} \tilde{\mathbf{R}} = \text{Tr}((\mathbf{M}^{-1})^T \mathbf{R} \mathbf{M}^T) = \text{Tr} \mathbf{R} = M_{eff} \quad (55)$$

The invariance of the regularization parameter w follows from the invariance of its defining relation (9). With $\Delta \mathbf{a} = (\mathbf{a} - \mathbf{a}^{obs})$ one verifies

$$\begin{aligned} \widetilde{\Delta \chi^2} &= (\widetilde{\Delta \mathbf{a}})^T \tilde{\mathbf{R}} \tilde{\mathbf{C}}^{-1} \tilde{\mathbf{R}}^T \widetilde{\Delta \mathbf{a}} \\ &= (\Delta \mathbf{a})^T \mathbf{M}^T (\mathbf{M}^{-1})^T \mathbf{R} \mathbf{M}^T (\mathbf{M} \mathbf{C} \mathbf{M}^T)^{-1} ((\mathbf{M}^{-1})^T \mathbf{R} \mathbf{M}^T)^T \mathbf{M} \Delta \mathbf{a} \\ &= (\Delta \mathbf{a})^T \mathbf{R} \mathbf{C}^{-1} \mathbf{R}^T \Delta \mathbf{a} \\ &= \Delta \chi^2. \end{aligned} \quad (56)$$

References

- [1] V. Blobel, “Unfolding Methods in High Energy Physics Experiments”, DESY 84-118.
- [2] J.E. Shore and R.W. Johnson, IEEE Trans. Inf. Theory IT-26 (1980) 26;
J.E. Shore and R.W. Johnson, IEEE Trans. Inf. Theory IT-29 (1983) 942.
- [3] Y. Tikochinsky, N.Z. Tishby and R.D. Levine, Phys. Rev. Lett. 52 (1984) 1357.
- [4] S. Kullback, “Information Theory and Statistics”, Wiley, New York (1959).
- [5] C.F. Gauss, “Theoria Motus Corporum Coelestium in Sectionibus Conicis Solem Ambientium”, Hamburgi Sumtibus Frid. Perthes et H. Besser Liber II, Sectio III (1809);
C.F. Gauss, “Theoria Combinationis Observationum Erroribus Minimis Obnoxiae”,
pars prior (15.2.1821) et pars posterior (2.2.1823),
Commentationes Societatis Regiae Scientiarum Gottingensis Receptiores
Vol. V (MDCCCXXIII);
- [6] M. Schmelling, Ph.D. Thesis, Dortmund (1987) ;
SFM Collaboration, A. Breakstone et al., Z. Phys. C40 (1988) 207 .
- [7] V.P. Zhigunov, Nucl. Instrum. and Methods 216 (1983) 183.
- [8] C.E. Shannon, Bell Systems Tech. J. 27 (1948) 379.
- [9] E.T. Jaynes, Phys. Rev. 106 (1957) 620.
- [10] Y. Alhassid, N. Agmon and R.D. Levine, Chem. Phys. Lett. 53 (1978) 22;
Y. Alhassid, N. Agmon and R.D. Levine, J. Comp. Phys. 30 (1979) 250.

Figure 1: Numerical unfolding example with 20 bins both for the true and the observed distribution. Figure (a) shows true, expected and actually observed distribution with its statistical errors. Figure (b) compares the true distribution to the unfolding result obtained by the MRX. The inset shows the result of unfolding by simply inverting the response matrix. A detailed discussion is given in section 4.

Figure 2: Limiting cases of the MRX. In figure (a) the MRX estimate for an underconstrained problem is compared to the result obtained according to the principle of minimum cross-entropy, figure (b) compares for the overconstrained case the MRX result to the least squares estimate.

Figure 3: Illustration of the criterion used to determine the regularization parameter w for the example given in figure 1. The width of the band shows the variation due to the measurement errors. For a detailed discussion see appendix D.