Astroparticle Physics 73 (2016) 44–51

Contents lists available at ScienceDirect

Astroparticle Physics

journal homepage: www.elsevier.com/locate/astropartphys

A likelihood method to cross-calibrate air-shower detectors

Hans Peter Dembinski [a,b,*], Balázs Kégl [c], Ioana C. Mariş [d], Markus Roth [a], Darko Veberič [a,c]

[a] Institute of Nuclear Physics, Karlsruhe Institute of Technology, Karlsruhe, Germany
[b] Bartol Institute, University of Delaware, Newark, DE 19716, United States
[c] Laboratoire de l'Accélérateur Linéaire, Orsay, France
[d] Universidad de Granada, Granda, Spain

ARTICLE INFO

Article history:
Received 2 May 2015
Revised 23 July 2015
Accepted 15 August 2015
Available online 20 August 2015

Keywords:
Air shower
Calibration
Pierre Auger Observatory
Fluorescence detector
Surface detector array

ABSTRACT

We present a detailed statistical treatment of the energy calibration of hybrid air-shower detectors, which combine a surface detector array and a fluorescence detector, to obtain an unbiased estimate of the calibration curve. The special features of calibration data from air showers prevent unbiased results, if a standard least-squares fit is applied to the problem. We develop a general maximum-likelihood approach, based on the detailed statistical model, to solve the problem. Our approach was developed for the Pierre Auger Observatory, but the applied principles are general and can be transferred to other air-shower experiments, even to the cross-calibration of other observables. Since our general likelihood function is expensive to compute, we derive two approximations with significantly smaller computational cost. In the recent years both have been used to calibrate data of the Pierre Auger Observatory. We demonstrate that these approximations introduce negligible bias when they are applied to simulated toy experiments, which mimic realistic experimental conditions.

1. Introduction

The latest generation of air-shower detectors, the Pierre Auger Observatory [1,2] and Telescope Array [3], are hybrid instruments. They combine a fluorescence detector which measures the calorimetric energy of an air shower with a low duty cycle, and a surface detector array measuring its size at ground with a full duty cycle.

The size of an air shower measured at the same point in its longitudinal development is proportional to a power of its energy [4]. Therefore, a calibration function returning an energy estimate for a measured size can be found by analyzing a subset of coincident events recorded in both detectors.

Fitting the calibration function to pairs of energy and size estimates with a plain least-squares method yields biased results for several reasons. Firstly, the least-squares approach requires the true energy of the air shower to be known event-by-event, but the fluorescence detector only provides an energy estimate that fluctuates around the true energy. Secondly, the energy spectrum of cosmic rays is very steep so that most of the data are located near the lower energy threshold of the detector.

In the threshold region, the detector triggers are not fully efficient. Upward fluctuations have a higher chance of passing the trigger and entering the data set than downward fluctuations. This creates an acceptance bias, so that the mean size of the accepted events does not reflect the true mean size of the original sample.

Applying an energy cut with a minimum energy high enough to avoid the threshold region altogether solves this problem, but it creates a new bias, caused by event migration over the new threshold introduced by the cut. How the bias appears is illustrated in Fig. 1. A superficial solution is to use a slanted cut, but determining the angle under realistic conditions, where the resolutions vary with energy and size of the air shower, requires Monte-Carlo simulation of the data [5].

We will show that a probabilistic approach solves the problem in a consistent way. Based on the known properties of air-shower development and the detectors, we construct a probability density model for the experimental data. Maximizing the likelihood of the data under this model then yields an asymptotically unbiased estimate of the calibration curve.

2. Definition of variables

We use the variable S for the size of the air shower at the ground, where it is observed by surface detector arrays. The size S depends on the energy E of the air shower, mass A, and geometry $\mathbf{a}$. We use air-shower geometry as a general term for the orientation and impact point of the air-shower axis. The size S is often obtained by fitting an empirical lateral distribution function to the ground signals [6,7],

* Corresponding author. Tel.: +1 302 8970079.
E-mail address: hdembins@udel.edu (H.P. Dembinski).

http://dx.doi.org/10.1016/j.astropartphys.2015.08.001
0927-6505/

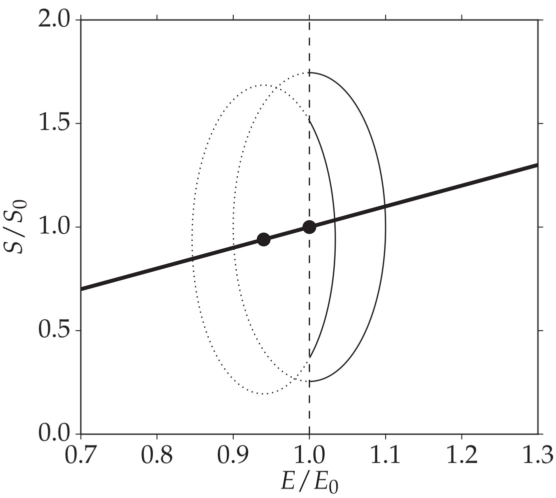

Fig. 1. Sketch of the bias introduced by an energy cut in a least-squares fit. Shown is an ideal calibration curve $S \propto E$ of the size S of an air shower against its energy E (thick solid line). Both are given in arbitrary units E_0 and S_0. An energy cut is placed at $E = E_0$ (vertical dashed line). Measured estimates of E and S fluctuate around the true values and spread events from points on the ideal calibration curve (black dots) outwards. We regard the simplest case of uncorrelated Gaussian smearing, where the iso-density contours of points scattered in this way form overlapping ellipses (only two are shown). Events that migrate below the cut line, which affects the measured estimates, are discarded (dotted part of ellipses). Surviving events (solid part of ellipses) fluctuate more often below the ideal calibration curve since the arc is longer. A least-squares fit wrongly compensates this by placing the calibration curve below the true curve.

but other proxies work as well, such as the inferred total number of muons at ground in very inclined showers [8].

Air showers with the same geometry $\mathbf{a}$ and energy E show a fluctuating size S at the ground. These fluctuations [9–11] are caused by random outcomes of the first few interactions of the air-shower development and possibly from sampling a random mass A from the mass-distribution of cosmic rays. The mass A is never exactly known event-by-event and therefore the dependency $S(A)$ adds to the observed fluctuations of S. We call these fluctuations combined *intrinsic fluctuations*.

Our aim is to find the function that yields the mean size $\bar{S}$ of the air shower, averaged over intrinsic fluctuations, as a function of its energy E and geometry $\mathbf{a}$. The energy dependence is usually modeled well by a power law $p_0 E^{p_1}$. Our approach does not depend on the exact relationship and therefore we will just refer to $\mathbf{p} = (p_0, p_1, \ldots, p_n)$ as the parameter vector of the function $\bar{S}(E, \mathbf{a}, \mathbf{p})$ that describes the energy dependence.

We mention the dependence of $\bar{S}$ on the full air-shower geometry $\mathbf{a}$ to treat the most general case. In practice, the dependency on $\mathbf{a}$ is usually corrected before applying the energy calibration. The correction is either based on air-shower simulations [7,8], or inferred from data, by demanding that the flux of cosmic rays looks isotropic in the corrected size [12].

Inverting $\bar{S}(E)$ gives the energy calibration function, which provides an energy estimate E_S based on a size S of the air shower. Care must be taken, however, since the random fluctuations of the observed size propagate into the energy estimate. Analyses based on E_S need to take into account, that E_S randomly fluctuates around the true energy E event-by-event, combined with the fact that true energies follow a very steeply falling distribution. This makes it more likely that a particular observed value of E_S was generated by an upward fluctuation of an air-shower of lower energy, than by one with the same or higher energy. If the distribution of energies E is to be measured based on E_S [6], unfolding methods can be used [14,15].

In addition to the effects discussed before, detectors do not measure the energy E, size S, and geometry $\mathbf{a}$ of the air shower directly. They provide estimates $\hat{E}$, $\hat{S}$, and $\hat{\mathbf{a}}$, that randomly fluctuate around the true values. These fluctuations are caused by statistical sampling of air-shower particles in the detector and by variations in the detector response. An experiment therefore provides a sample of tuples $(\hat{E}_i, \hat{S}_i, \hat{\mathbf{a}}_i)$ as input for the analysis. We assume that an energy cut $\hat{E} > E_{\text{cut}}$ is applied to this set which discards events with poor resolution in the threshold region of the detector.

To distinguish between functions and probability density functions (pdfs) in this article, we use the semi-colon in pdfs to separate the random variables from the dependent variables. For example, $f(x; p)$ is the probability density function f of the random variable x, whose location and shape depends on p. When integrals over random variables appear, we will not explicitly indicate the limits, except if the integral does *not* cover the physical domain of the variable, for example, $[0, \infty)$ for E and S.

We will refer to the normal distributions frequently, and therefore use the notation $\mathcal{N}(x; \mu, \sigma)$ to indicate the density

$$\mathcal{N}(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2\right). \tag{1}$$

In a fully rigorous treatment, we would have to use the truncated normal distribution in most places, where the domain of the variable x is not the full real line. We generally assume that the experimental conditions are such that $\mu/\sigma \gg 0$, so that both distributions approach each other.

3. Likelihood estimation of the calibration function

Our fitting method is based on the maximum-likelihood method [13]. For un-binned continuous data, it states that an estimate of the parameter vector $\mathbf{p}$ can be found by maximizing the joint pdf $\mathcal{L}$ of the data under the model considered. We make a usual substitution and maximize $\ln \mathcal{L}$ instead of $\mathcal{L}$,

$$\ln \mathcal{L}(\mathbf{p}) = \sum_i \ln f(\hat{E}_i, \hat{S}_i, \hat{\mathbf{a}}_i; \mathbf{p}), \tag{2}$$

which is equivalent but easier to handle. The density $f(\hat{E}, \hat{S}, \hat{\mathbf{a}}; \mathbf{p})$ models the data distribution as a function of $\mathbf{p}$. We maximize this sum with standard numerical algorithms to get an estimate $\hat{\mathbf{p}}$ of $\mathbf{p}$.

If the data density was very high, working with a histogram of the data would be more effective and the log-likelihood would take a different form. Both approaches can also be combined, so that the former is used in high density regions to speed up the computation of the sum, an example of such a technique is given in Ref. [15].

The maximum-likelihood approach has a useful property that we will exploit repeatedly. Finding the maximum of $\ln \mathcal{L}$ to get the estimate $\hat{\mathbf{p}}$ only involves the first derivative $\nabla_{\mathbf{p}} \ln \mathcal{L}$. Similarly, computing the uncertainty estimate of $\hat{\mathbf{p}}$ only involves the second derivative. Therefore, any constant factors c_i with $\nabla_{\mathbf{p}} c_i = 0$, that appear in the evaluation of $f_i(\mathbf{p}) = f(\hat{E}_i, \hat{S}_i, \hat{\mathbf{a}}_i; \mathbf{p})$, can be dropped without changing these results,

$$\begin{aligned} \ln \mathcal{L}(\mathbf{p}) &= \sum_i \ln f_i(\mathbf{p}) = \sum_i \ln c_i\, f_i'(\mathbf{p}) \\ &= \sum_i \ln c_i + \sum_i \ln f_i'(\mathbf{p}) \equiv \sum_i \ln f_i'(\mathbf{p}). \end{aligned} \tag{3}$$

We will use this to avoid the explicit computation of such factors wherever possible.

We now focus on the construction of $f(\hat{E}, \hat{S}, \hat{\mathbf{a}}; \mathbf{p})$. The size function $\bar{S}(E, \mathbf{a}, \mathbf{p})$ of the air-shower is at the heart of this pdf, the crucial point is to model the random fluctuations of events around this mean.