

Maximum-likelihood detection of sources among Poissonian noise

I. M. Stewart

Jodrell Bank Centre for Astrophysics, University of Manchester, Oxford Road, Manchester M13 9PL, UK
e-mail: Ian.Stewart-2@manchester.ac.uk

Received 7 November 2008 / Accepted 31 December 2008

ABSTRACT

A maximum likelihood (ML) technique for detecting compact sources in images of the X-ray sky is examined. Such images, in the relatively low exposure regime accessible to present X-ray observatories, exhibit Poissonian noise at background flux levels. A variety of source detection methods are compared via Monte Carlo, and the ML detection method is shown to compare favourably with the optimized-linear-filter (OLF) method when applied to a single image. Where detection proceeds in parallel on several images made in different energy bands, the ML method is shown to have some practical advantages which make it superior to the OLF method. Some criticisms of ML are discussed. Finally, a practical method of estimating the sensitivity of ML detection is presented, and is shown to be also applicable to sliding-box source detection.

Key words. methods: data analysis – techniques: image processing – X-rays: general

1. Introduction

X-ray observatories from HEAO-2/Einstein onward have carried imaging cameras which have counted individual X-ray photons as they impinged upon a rectilinear grid of detector elements. The number of X-ray events detected per element within a given time interval will be a random integer which follows a Poisson probability distribution.

As the sensitivity of detectors and the effective areas of X-ray telescopes have increased, so has the X-ray cosmic background been increasingly resolved into compact sources. E.g. a telescope such as XMM-Newton detects on average 70 serendipitous sources within the field of view of each pointed observation (Watson et al. 2008). A full characterisation of the X-ray background would shed light upon the evolution of the early universe (see Brandt and Hasinger 2005, for a recent review of deep extragalactic X-ray surveys and their relation to cosmology). It is therefore of interest to find the best way to detect compact sources in these images – that is, the detection method which offers the best discrimination between real sources and random fluctuations of the background.

In a previous paper (Stewart 2006, hereinafter called Paper I) I discussed the general class of linear detection methods, i.e. methods which form a weighted sum of the values of several adjacent image pixels, and optionally also over several different energy bands. A method for optimizing the weights was presented and shown to be substantially superior to the sliding-box method which, no doubt because of its simplicity, has been frequently used for compiling catalogs of serendipitous sources (see e.g. DePonte & Primini 1993; Dobrzycki et al. 2000; also task documentation for the exsas task *detect* and the SAS task *eboxdetect*).

The present paper should be considered as an extension to Paper I and also partly as an emendation. Several issues discussed in Paper I either had not at that time been completely

thought through, or were not described as clearly as they might have been. These issues are itemized as follows:

- the difference between the distribution of samples of a random function, and the distribution of its peak values, was not realized;
- the concept of sensitivity was not adequately defined;
- the normalization of the signal did not allow direct comparison between single-band and multi-band detection sensitivities;
- both Paper I and the present paper are concerned not only with numerical values of the sensitivity of various methods of source detection, but also with practical, approximate methods of calculating these values “in the field”. One may prefer a lengthy, complicated method for calculating values for a paper; such is unlikely to be an equally satisfactory field method. In Paper I there was insufficient grasp of this conceptual distinction;
- sensitivities for the *N*-band *eboxdetect* method were incorrectly calculated (see Sect. 4.2 for an explanation).

The different signal normalization adopted in the present paper precludes direct comparison of the results herein with those of Paper I; however after appropriate rescaling the two sets of results are consistent. No numerical errors in Paper I were found, despite extensive reworking of the software for the present paper.

The primary aim of the present paper is to present a method of source detection based on the Cash statistic (Cash 1979) and to show that this is at least as sensitive as the optimized linear filter described in Paper I. For this purpose, sensitivities of a variety of detection methods are compared across a range of background values for both single-band and multi-band detection scenarios. The paper also aims to demonstrate other advantages of the Cash source detection, and to discuss and investigate some reported disadvantages.

Section 2.1 contains a mathematical description of the general source detection problem. In Sect. 2.2 a 3-part detection

algorithm is described, in which a map of the detection probability at each image pixel is made as step one, peaks (maxima) in this map are located in step two, and a source profile is fitted to each peak in step three. The relation between the respective probability distributions of the map vs. peak values is also discussed. Detection sensitivity is defined in Sect. 2.3.

Cash-statistic source detection is described in Sect. 2.4.1. In Sect. 2.4.2 it is shown how the Cash statistic may be used at both stages 1 and 3 of the 3-part source detection algorithm. Finally, in Sects. 2.4.3 and 2.4.4, objections to the Cash method are addressed.

Section 3 contains the results of the sensitivity comparisons. In Sects. 3.1 and 3.2, details of the necessary calculations are discussed. 1-band results are presented in Sect. 3.3 and N -band results in Sect. 3.4. The degradation of some N -band sensitivities when the true source spectrum does not match that assumed is discussed in Sect. 3.5.

An approximation useful for “in the field” estimation of Cash sensitivities is presented in Sect. 4.1. Finally, in Sect. 4.2, a deficiency of the sliding-box method is discussed, and a method for remedying it is proposed.

2. Theory

2.1. Source detection

In the situation relevant to the present paper we have a model function $e(\mathbf{r})$ defined over the space \mathbf{r} by

$$e(\mathbf{r}) = B(\mathbf{r}) + \sum_j^M \alpha_j S_j(\mathbf{r} - \mathbf{r}_{0,j})$$

where B is called the background, each α is a scalar amplitude and S is the signal shape. We conceive of a set of N positions \mathbf{r}_i for $i \in [1, N]$, each with an associated measurement of flux c_i . Each c_i is a random integer with a Poisson distribution about a mean given by the value e_i of the model at that position. From this comes the requirement that $e_i \geq 0 \forall i$. In mathematical shorthand

$$\langle c_i \rangle = e_i = B_i + \sum_j^M \alpha_j S_j(\mathbf{r}_i - \mathbf{r}_{0,j}).$$

In CCD X-ray detectors the c_i are obtained by accumulation within arrays of voxels, i.e. volumes within the physical coordinates \mathbf{r} . In this case the functions B and S are integrated over the i th voxel dimensions to give the respective expectation values $B_i, S_{i,j}$ etc. For the purposes of the present paper it is mostly assumed that the coordinates \mathbf{r} consist of just the spatial (x, y) coordinates of the X-ray detector; in some sections the energy of the X-ray photons is added to the set. In this situation each voxel is bounded in the (x, y) directions of course by the edges of the physical pixels of the detector. In the energy direction, voxels are bounded by the boundaries of any energy selection performed by the user.

The ultimate aim of source detection is to obtain the best estimates possible of α and \mathbf{r}_0 for the most numerous possible subset of the M sources present. Satisfactory performance of this depends upon a couple of conditions. Firstly, the sources must not be confused – i.e. the density of “detectable” sources should be $\ll 1/\Delta\mathbf{r}_S$ where $\Delta\mathbf{r}_S$ is some characteristic size of S . For XMM-Newton this requires the detectable source density to be less than about 10^5 deg^{-2} . The deepest surveys to date reach X-ray source densities of only a tenth of this value

(Brandt & Hasinger 2005). On the other hand, it was shown in Paper I that perturbations of the background statistics set in at much shallower levels of sensitivity; but still a little beyond the deepest practical reach of XMM-Newton. There appears thus to be still some scope to improve source detection in the present observatories without coming up against the confusion limit. To some extent also one can disentangle confused sources during the fitting process, by comparing goodness-of-fit for fitting one versus two or more separate source profiles to a single detection.

If we may neglect confusion then we can detect sources one at a time. The model formula can thus be simplified to

$$e_i = B_i + \alpha S(\mathbf{r}_i - \mathbf{r}_0). \quad (1)$$

The second condition necessary to source detection is that the form of S is (i) known (at least approximately); and (ii) not too similar to B . In X-ray detectors used to date, a significant fraction of sources have angular sizes which are smaller than the resolution limit of the telescope. S for such sources is then simply the PSF of the telescope, which can be estimated a priori, and which is often much more compact, i.e. with more rapid spatial variation, than the background. However, about 8 percent of sources detected by XMM-Newton for example are resolved or extended (Watson et al. 2008). Detection of a resolved source is much more difficult (or at least, less efficient) since its S is dominated by the angular distribution of source flux, which is not known a priori. In addition, there is of course no upper limit to the size of X-ray sources, and in fact the distinction between source and background is largely a philosophical one. In practice one has to impose an essentially arbitrary size cutoff and consider only variations in X-ray flux which have spatial scales smaller than that cutoff to be sources, and take everything else to be background.

Modern CCD detectors can measure not only the positions of incident X-ray photons but also their energies. This opens an additional degree of freedom for S and B and thus the possibility of exploiting differences in their respective spectra to separate them with even greater efficiency. The average spectrum of the sources in the 1XMM catalog is quite different from the instrumental background spectrum of typical instruments (see Carter & Read 2007; and Kuntz & Snowden 2008, for recent data on XMM-Newton), and appears also to be different to the spectrum of the cosmic background (see e.g. Gilli et al. 2007). However if it is true, as most authorities now suppose, that practically the whole of the cosmic X-ray background at energies higher than a few keV is comprised of point sources, one might expect the latter difference to diminish as fainter sources become detectable.

Parallel detection of sources in images in several energy bands of the same piece of sky is complicated by the fact that, although all (point) sources have the same PSF at a given detector position and X-ray energy, there is no corresponding common source spectrum. This means that we cannot “locate” the source S in the energy dimension – we must be satisfied with locating it spatially on the detector.

Some of the detection algorithms described in the present paper avoid making any assumption about the source spectrum; others require the user to choose a spectrum a priori. As will be seen from the results, there is usually a trade-off involved, in that methods which require the user to select a spectrum template may be better at detecting sources with spectra close to that template, but worse where the source spectrum is very different to the template; whereas a method which makes no assumptions will usually perform with about the same efficiency on all sources.

2.2. Likelihood-map source detection

The present paper treats of a generic source-detection procedure which falls into three parts. The first part calculates a map giving an initial, possibly coarse estimate of the probability, for each pixel, that a source is centred on that pixel. In other words, it tests, for each j th map pixel, the following flux model:

$$e_i = B_i + \alpha S(\mathbf{r}_i - \mathbf{r}_j). \quad (2)$$

This task is accomplished by calculating some statistic U from the count measurements c_i taken from some subset of the N detector pixels. The spatial distribution of this subset ought not to be much larger than the size of S – otherwise one is including pixels which one knows contain no information about S . In other words, the pixels chosen to calculate the statistic at detector pixel j are best located within a field of approximate size Δr_S , centred about pixel j .

After calculating the statistic U for all pixels, this first part of the detection procedure converts this into a measure of the probability that the value of U at pixel j would arise by chance fluctuation of the background alone.

The second part of the source-detection procedure consists of a peak detection algorithm. The output from this part is a list of peaks in the probability map, which are interpreted as source candidates.

The final task cycles over this list of candidates and attempts to fit the signal profile to the field of c_i values surrounding each. To be of maximum use, the fitting procedure should be capable both of separating a single candidate into two or more confused sources, and merging multiple candidates which turn out to belong to a single source.

In the XMM-Newton data-product pipeline, both stages 1 and 2 are performed by the SAS task called *eboxdetect*, whereas the final stage is performed by *emldetect*.

The final task may perform some additional calculations. It may calculate χ^2 for the fit so as to aid in discriminating real cosmic sources from other phenomena, such as “hot” CCD pixels, which may give rise to excess counts (this is not at present done by *emldetect*). As is shown in Sect. 2.4, use of the Cash statistic allows one to calculate a final value for the source-detection likelihood after the fit.

This separation of detection and fitting stages is a useful procedure only where the image pixels are not significantly larger than the characteristic width Δr_S of S . This criterion is in fact achievable with the imaging cameras of all three current observatories primarily dedicated to detecting X-rays, namely XMM-Newton, Suzaku and Chandra (Strüder et al. 2001; and Turner et al. 2001, for XMM-Newton; Serlemitsos et al. 2007; and Koyama et al. 2007, for Suzaku; and Weisskopf et al. 2002, for Chandra).

2.2.1. Peak detection

The peak-detection stage does not seem to have been the subject of much study. It is problematical for two reasons. Firstly, the frequency distribution of values at the maxima of a random function (call it the “peak” distribution) is not in general the same as the distribution of values over a set of randomly chosen positions (which we might call the “whole-map” distribution). In general one would expect the former distribution to be somewhat broader than the latter, since the probability that a random pixel is a peak is small for relatively low values but large for high values. Nakamoto et al. (2000) give a method to calculate

the distribution of peaks for a random function of one dimension, but it is not clear if this theory can be extended to more than one dimension. However, as is shown in Sect. 2.4.4, the Cash maximum likelihood statistic does provide a way to estimate the shape of the peak distribution, although not its normalization (which depends on the fraction of map pixels which are peaks). The normalization must be estimated by appropriate simulations.

The second problem is that the definition of a peak in an array of samples of a nominally continuous function is somewhat arbitrary. One obvious choice of definition is to label any j th pixel as a peak pixel if the value c_j at that pixel is greater than the values at all other pixels within a patch of pixels surrounding the j th. But what size and shape to make the patch? It seems clear that these ought to match the size and shape of S – too large a patch might miss additional, nearby sources; whereas too small runs into the opposite danger of listing more than one peak for a single source. Also, it is not clear how this choice might affect the frequency distribution of peak values. This danger was pointed out by Mattox et al. (1996) in respect of source detection in EGRET data.

2.3. Sensitivity

The efficiency of a method of detecting sources is characterised by its sensitivity, which, broadly speaking, is the smallest amplitude α which will result in a source signal αS being detected among background of a given level. While this sentence conveys the general idea, a precise mathematical definition is necessary before a comparison can be done between methods of source detection. A somewhat stilted and obscure definition of sensitivity was presented in Paper I. The present section is an effort to make this definition clearer and more rigorous.

2.3.1. Definition

Let us start by assuming that B and S are known a priori. Formally speaking this is not true: one ought to make use only of their best estimates \hat{B} and \hat{S} in the formula for any source detection statistic. In practice however, because B and S can often be quite well estimated by methods separate from the source detection itself, it is convenient for the sake of simplicity to make the approximation $\hat{B} \sim B$ etc.

Each of the detection methods discussed in either Paper I or in the present paper can be associated with a statistic U which is to be evaluated at positions of interest: perhaps at the centre of each image pixel, or perhaps at a set of positions of candidate sources. U at any position is evaluated, using a formula particular to the detection method, from the measured counts c_i . Clearly U so calculated is a random variable which will occur with some well-defined probability distribution $p(U)$. We can also define an integrated (more strictly, a reverse-cumulative) probability P such that

$$P(>U) = \int_U^{\infty} dU p(U).$$

In order to understand how sensitivity relates to U it is helpful to imagine an ensemble of measurements of a source of a given amplitude α . Or one could simulate this situation via a Monte Carlo experiment which, for each member of its ensemble, generates a set of c_i values. The input model for this should consist of Eq. (2) and can make use of the a priori known shapes of B and S . For each member of the ensemble, U is calculated from the c_i values. The resulting ensemble of U values will have a distribution

which will in general be different for different a priori-assumed values of α . To reflect this, let us write both the probability density p and the reverse-cumulative probability $P(>U)$ with an additional α argument as $p(U; \alpha)$ and $P(>U; \alpha)$.

We are interested first of all in the predicted P in the case that α equals zero – i.e., when there is no source in the detection field, only background. It is necessary to consider this scenario because the random nature of the data means that, at least some of the time, we will unavoidably be labelling as a source a concentration of counts which is merely a random fluctuation of the background.

The first step in source detection is to decide what fraction of these false positives we are prepared to put up with. This decision defines a cutoff probability P_{det} which in turn is associated (through the function $P(>U; 0)$) with a definite value of U_{det} .

As an aside, note that in all cases in which the measured data are random integers, $p(U; \alpha)$ is a sum of delta functions, thus $P(>U; 0)$ is a descending step function. Inversion of the latter to obtain U_{det} can clearly only give an approximation in which U_{det} is set to the next highest value of U for which p is defined. In this situation one may expect the sensitivity of the method to descend in steps with decreasing background. This effect can be seen in the plots of the sensitivity of the simple sliding-box detection scheme in Figs. 5, 7 and 8.

The sensitivity α_{det} of a detection method is here defined to be that value of α which, for given B and S , gives rise to values of U such that $\langle U \rangle = U_{\text{det}}$, where U_{det} is obtained from P_{det} as described above. See Fig. 1 for a diagrammatic example of this. The algorithm for calculating sensitivity is thus as follows:

1. decide a value for P_{det} ;
2. invert the relation $P = P(>U; 0)$ (solid line in Fig. 1) to obtain U_{det} from P_{det} ;
3. an expression for the average value of U as a function of α is inverted to give $\alpha_{\text{det}} = \langle U \rangle^{-1}(U_{\text{det}})$.

Note that other definitions are possible, for example one could define α_{det} such that the median of $U(\alpha_{\text{det}})$ instead of its mean was adjusted so as to equal U_{det} . I have preferred the definition in terms of the mean simply because it is generally easier to calculate. However in the limit of high counts, all the varieties of U here studied tend to a gaussian distribution, for which the mean and median are equal. Within the range of background values treated in the two papers, the median and mean definitions empirically are found to produce similar values.

The mean definition is equivalent to the “counts amplitude” concept employed in Paper I; however the median definition is implicit in Fig. 12 of that paper.

The curves in Fig. 1 were constructed via Monte Carlo. At each iteration, simulated data values were generated for a small (9×9) patch of pixels; the value of the appropriate detection statistic U was then calculated for the patch. The end result of the Monte Carlo is an ensemble of U values. The data values were Poisson-distributed integers whose mean values were given by a model comprising flat background of 1.0 counts/pixel plus a centred source profile. The solid curve and the error bars represent Monte Carlo-generated ensembles with 10^6 and 10^5 members respectively.

As mentioned in Sect. 2.2, a multi-band detection method which does not require the user to make an a priori choice of source spectrum is expected to perform better at “across the board” detection, because of the fact that cosmic X-ray sources don’t adhere to a common spectral template. Note however that, regardless of the requirements of a (multi-band) detection

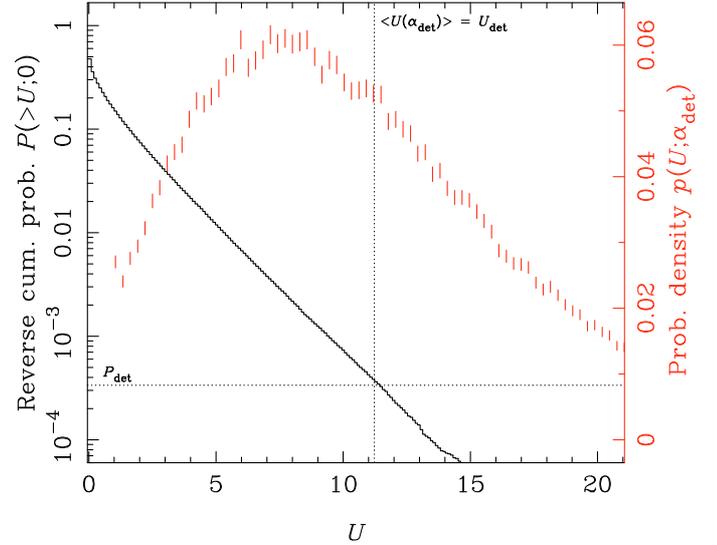


Fig. 1. The solid curve and the error bars (the latter shown in red in the electronic version) represent the results of separate Monte Carlo experiments, as described in the text. The solid curve shows the reverse-cumulative probability $P(>U; \alpha)$ in the case that the model consists entirely of background (i.e., $\alpha = 0$). The horizontal dotted line shows the chosen detection probability P_{det} . Its interception with the solid curve gives the cutoff value U_{det} of the detection statistic U . The error bars show the approximate course of the probability density function $p(U)$ for a case in which the model contains a non-zero source component. In the present case the model source amplitude α has been adjusted such that $\langle U(\alpha_{\text{det}}) \rangle = U_{\text{det}}$. This is the required condition for α to equal the sensitivity α_{det} .

method, calculation of its sensitivity always requires the user to choose a spectral template. E.g. in the Monte Carlo thought experiment described above, if α is non-zero, S needs to be completely defined over the whole space \mathbf{r} , which in the multi-band case includes an energy dimension, in order to generate c_i values and ultimately a value for $\langle U \rangle$.

2.3.2. Map vs. peak sensitivity

In Paper I it was implicitly assumed that if one made a frequency histogram $p_{\text{peak}}(U)$ of U values at an ensemble of locations of detected sources, that this would be related to the probability distribution $p_{\text{map}}(U)$ of U values at all image pixels by the following proportional relation:

$$p_{\text{peak}}(U) = \frac{A_{\text{image}}}{A_{\text{beam}}} p_{\text{map}}(U)$$

where A_{image} is the image area in pixels, and A_{beam} represents a constant known as the “beam area”, taken to be approximately equal to the FWHM area in pixels of the PSF. If this were true, the cutoff probability P_{det} could be obtained by multiplying the maximum acceptable number of false detections per unit image area by the beam area. I now realize however that this assumption is not true: as described in Sect. 2.2.1 and demonstrated in Fig. 3, the shapes of the two distributions are different. It follows then that one ought to use the probability distribution of U values at peaks in the map, rather than at all map pixels, in calculating sensitivity.

This is a pity, because there are several attractive features of calculating “map” rather than “peak” sensitivities. Three of these have present relevance. The first and most obvious of these “attractions” is the retention of consistency with Paper I.

Secondly, it is usually much easier and more straightforward to calculate “map” sensitivities. An adequate closed-form approximation for the “map” version of the cumulative probability function P exists for at least some of the statistics considered herein, whereas the “peak” versions of P must, for all but the Cash statistic, be estimated via a lengthy Monte Carlo. In addition, calculation of “peak” distributions naturally also requires the addition of a peak detector to the algorithm; the Monte Carlo must create a large image rather than a small patch of pixels; boundaries of same have to be avoided, and so forth.

The third thing which makes map rather than peak calculations attractive arises because the true number density of false detections (indeed of detections of any kind) depends not just on the formula for U employed, but also on the nature of the peak-detection algorithm. Having to consider the effect of different algorithms within a “space” with two degrees of freedom rather than one is an unwelcome complication.

None of these nice features of the “map” calculation would justify its use, however, if it were not for the fact, as pointed out by Nakamoto et al. (2000), that the “map” and “peak” distributions of a random function become asymptotically similar towards high values of the function. This convergence is observed for example in Fig. 3. Since source detection by its nature deals with extreme values of the null-hypothesis probability function, it was felt to be acceptable to retain the “map” calculation for Sect. 3, in which sensitivities of the various methods are compared. However when it comes to calculation of the absolute sensitivity of the Cash (or any other) statistic, of course the true “peak” version of P must be employed.

2.4. The Cash statistic

2.4.1. Definition

Cash (1979) discussed a χ^2 -like statistic for Poisson data. Cash’s statistic was constructed in two stages. The first stage comprises the log-likelihood expression

$$L = \sum_{i=1}^N (c_i \ln[e_i] - e_i - \ln[c_i!])$$

where e_i is a model function for the expectation value $\langle c_i \rangle$. L can be seen to consist of the logarithm of the product of the individual Poisson probability densities $p(c_i | e_i)$. If the model e can be expressed in terms of one or more adjustable parameters, we can define the best fit values of these parameters as those which produce the largest value of L . Note that the factorial term is not a function of e and thus can be neglected for fitting purposes.

The second stage of Cash’s development was to form the difference

$$U_{\text{Cash}} = 2(\max[L_q] - L_{\text{null}}) \quad (3)$$

where L_{null} is L calculated by setting all parameters of the model e to fixed, background values and $\max(L_q)$ for integer q is shorthand for the maximum result obtainable by varying a pre-selected number q of parameters of e .

So long as the null hypothesis model can be expressed by a combination of legal values of the q free parameters, $U_{\text{Cash}} \geq 0$. According to Cash’s theory, under this and other conditions discussed below, U_{Cash} in the asymptotic case is distributed approximately as χ_q^2 , i.e. χ^2 for q degrees of freedom.

Just as with χ^2 for gaussian data, the Cash statistic can be used not only to fit parameter values of the model but also to obtain information about the probability distribution of those

values. Cash for example uses it to derive confidence intervals for fitted parameters. The XMM-Newton SAS task *emldetect* (online description at <http://xmm.vilspa.esa.es/sas/current/doc/emldetect.ps.gz>) employs the Cash statistic firstly to obtain best-fit values for the position, amplitude and optionally extent of sources in X-ray exposures, and secondly, at the conclusion of the fitting routine, to calculate a detection probability for each so-fitted source. Calculation of the detection probability is done by setting the source amplitude to zero in the expression for L_{null} . The probability of obtaining the resulting value of U_{Cash} is then equivalent to the probability of obtaining the fitted values of the source parameters when there is nothing in the field but background.

Descriptions of the use of the Cash statistic for source detection can be found in Boese & Doebereiner (2001) and Mattox et al. (1996) to give just two examples. The expressions developed in the present paper are not based on either treatment but comparison reveals close similarities.

2.4.2. Map vs. peak Cash

In the present paper, the Cash statistic is evaluated under two different scenarios: either at the centre of each image pixel to make a map of U_{Cash} samples, or at the best-fit locations of detected sources. In the first case, the source is taken to be centred on the pixel in question (Eq. (2)): the amplitude is the only fitted (i.e. free) parameter. In the second case, the x and y coordinates of the source centre are fitted as well as the amplitude. It is as well to consider how the U values returned by these respective situations are related.

In either case, practical considerations limit the number of image pixels included in the calculation of each value of U_{Cash} to a relatively small patch surrounding the pixel of interest. As can be seen from Fig. 6 however, the size of this patch has an asymptotically decreasing effect on the result of the calculation as this size becomes much larger than the characteristic size of the PSF. Consider then an ideal, continuous, two-dimensional function U_{ideal} obtained by calculating U_{Cash} , with only the model amplitude free, at all points within the image boundary. Samples of this smooth function on the grid of points at the centres of the image pixels represent the asymptotic values of the actual, practical map- U evaluation in the limit of large “patch” size. The values at peaks of U_{ideal} bear the same relationship to the actually calculated values of peak- U at the end of the final fitting procedure. The addition of two more fitted parameters does not change the value of U at these peaks: only the appropriate probability distribution.

These ideas can be expressed more compactly if we introduce a little notation. Let $U(A_{\text{map}})_{i,j}$ be the “map” value of U evaluated at the pixel centred on (x_i, y_i) , for a patch area A_{map} ; let $U(\hat{x}_0, \hat{y}_0, A_{\text{peak}})$ be the value at the fitted source location (\hat{x}_0, \hat{y}_0) , for patch area A_{peak} ; and let $U_{\text{ideal}}(x, y)$ be the ideal, smooth function. Then

$$U_{\text{ideal}}(x_i, y_i) = \lim_{A_{\text{map}} \rightarrow \infty} U(A_{\text{map}})_{i,j}$$

and

$$U_{\text{ideal}}(\hat{x}_0, \hat{y}_0) = \lim_{A_{\text{peak}} \rightarrow \infty} U(\hat{x}_0, \hat{y}_0, A_{\text{peak}}).$$

Imagine the 3-part detection scenario of Sect. 2.2 with the Cash prescription used both to calculate the map of U values in step 1 and to fit the source parameters in step 3. As mentioned above, an advantage of the Cash statistic is that it can then be used to

calculate a final value of U (thus of the null-hypothesis probability) for each fitted source. If the image pixels are smaller than the characteristic size of the PSF (mentioned in Sect. 2.2 as a condition for the usefulness of the 3-part algorithm), and if we assume (as seems reasonable) that the maxima in U_{ideal} are asymptotically paraboloid, then we expect two things: firstly, that the positions output by step 2, i.e. positions of maxima in the array of “map” values produced by step 1, will be, most of the time, within 0.5 pixels of the final fitted positions; secondly, and because of this, that the final U values will be approximately the same as, but slightly larger than, the U values at the step 2 grid positions. These results are borne out by Monte Carlo experiments.

The second expectation, plus the fact that, as mentioned in Sect. 2.3, the respective tails of the “map” and “peak” distributions of U are expected to be similar, together mean that we may investigate the sensitivity of Cash statistic source detection, or at least compare it with other methods, via the (much more convenient) “map” process alone.

Since this paper shows that the Cash statistic is at least as sensitive as, and more flexible than, an optimized linear filter, it might be supposed that a scheme of source detection which employed Cash for stage 1 as well as stage 3 would be optimum in practice. However, although it has been found convenient to make use of this scheme as described above in the present paper, it can be shown that this is probably not the best scheme for practical source detection. In fact it isn’t necessary to use the Cash statistic for step 1 as well as step 3; nor is this likely to be an optimum use of computer resources. An alternative method is well exemplified by the current XMM-Newton detection scheme, in which the detection statistic of step 1 is calculated using a sliding-box formula (see Eq. (B.1)). The sliding-box statistic is much quicker to calculate than the Cash statistic, and its relatively poor sensitivity is offset by setting the pass threshold low enough to pass any possible source through to step 3. Since step 3 does employ the Cash statistic, the full sensitivity of the latter is retained without a great penalty in computing time.

For purposes of the present paper, where the amplitude alone is fitted, this was done via a Newton-Raphson method, as detailed in Appendix A. Where more than one parameter was to be fitted, the Cash statistic was maximized via a Levenberg-Marquardt method (following the algorithm described by Press et al. 2003).

2.4.3. Objections to ML measures of significance

Various authors (e.g. Freeman et al. 1999; Protassov et al. 2002; Bergmann & Riisager 2002; Pilla et al. 2005; Lyons 2007) have criticized the use of Maximum Likelihood statistics for the detection of sources or spectral lines in Poisson-noise data. There seem to be two main concerns: firstly that the parameter values which return the null hypothesis lie on the boundary of the space of permissible values; secondly that the data to which a line (or source) profile is fitted are also those used to estimate the background (null hypothesis) signal.

The first objection appears to be a result of confusing physical with statistical constraints. Statistically speaking, for Poisson data, there is nothing wrong with a fit which returns a negative value for the amplitude, provided the net model flux $B + \alpha S$ remains non-negative. Even if one specifies that the physical model at issue does not allow negative amplitudes (although it is easy to conceive of situations which could give rise to a dip in flux at a particular position or energy), the fitted amplitude $\hat{\alpha}$

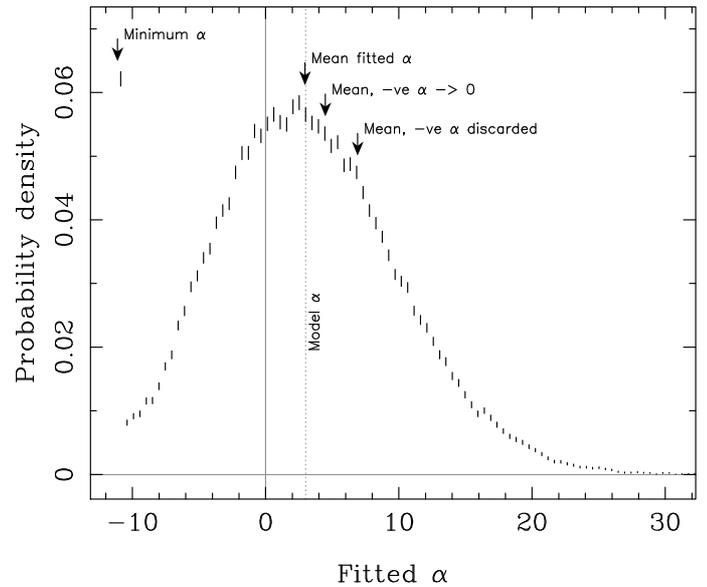


Fig. 2. The error bars are the result of a Monte Carlo experiment as described in the text. They show the frequency distribution of values of the source amplitude α , fitted by use of the Cash statistic. All members of the Monte Carlo were generated using the same model amplitude (3.0 counts), indicated in this figure by the vertical dotted line. Negative values of fitted α were allowed; but the fit was not permitted to return a value less than the statistical limit (indicated with an arrow and the label “Minimum α ”), which is dictated by the necessity for the total model flux to remain ≥ 0 for all pixels. The pile up of these bounded fits can be seen in the high numbers within the first bin. It can be clearly seen that it is necessary to retain the negative values of fitted α if the mean of the distribution is to be a good approximation to the model value of α .

and the model amplitude α are not the same thing: the former is only an estimator for the latter, thus is a random variable with a certain frequency distribution. Where the model amplitude is small compared to the scatter in the fitted amplitude, it is necessary to allow negative values in order for the fitted amplitude to be an accurate estimator for the model. Figure 2 demonstrates this. (Fig. 2 was constructed via Monte Carlo in a way similar to Fig. 1. The ensemble comprised 10^5 members.)

A negative value for $\hat{\alpha}$ which is small compared to its uncertainty σ_{α} , where the physical model specifies that $\alpha \geq 0$, merely indicates that the measurement is consistent with zero. In this case one could use the sensitivity of the statistic as an upper limit. If $\hat{\alpha}$ is negative but large compared to σ_{α} , the conclusion is rather that there is probably something wrong with the physical model (although, of course, such values will occur with a certain frequency by chance). In neither case are negative values of $\hat{\alpha}$ somehow incorrect. The only requirement which the statistical theory imposes in cases where the noise is Poissonian is that the total flux (background plus source) may not be negative. Thus one can conclude, in contradiction to Protassov et al. (2002), that the null value of α , i.e. zero, does not lie on the boundary of the permitted values of $\hat{\alpha}$.

Precisely what one does with negative $\hat{\alpha}$ values depends on one’s aim. If that aim is to obtain the best estimate of α , as in the above paragraphs, then negative $\hat{\alpha}$ s should be retained. However if one desires to compile a list of candidate sources, then fits giving negative $\hat{\alpha}$ ought probably to be discarded from the list. In no case is it correct to keep the candidate but tamper with the $\hat{\alpha}$ value (by, for example, setting it to zero).

It should be pointed out that policy toward negative $\hat{\alpha}$ values should differ depending on whether the detection is over a single

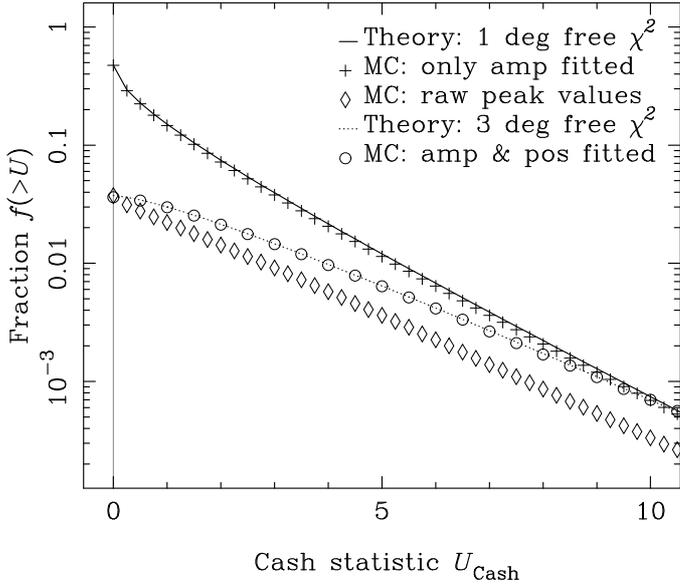


Fig. 3. Reverse-cumulative frequency histograms from Monte Carlo ensembles of the Cash statistic U_{Cash} . Continuous curves plot the theoretical distributions, which are just the χ_n^2 distributions for the respective n degrees of freedom, suitably normalized to the respective fraction of pixels involved. Crosses, circles and diamonds represent Monte Carlo data. The solid curve (theory) and crosses (Monte Carlo) show results where only the source amplitude is fitted (=1 degree of freedom), for a source centred at each image pixel. The diamonds give the distribution of U values for that small subset of these “map” pixels which were identified as peaks. Finally, the dotted curve (theory) and the circles (Monte Carlo) show the distribution of values obtained by fitting x and y coordinates as well as amplitude (=3 degrees of freedom) for each of the “map” peaks.

spectral band or over several in parallel. For single-band detection it is appropriate to discard candidates for which $\hat{\alpha}$ is negative – indeed, if a different algorithm is used for stage 1 of the detection process, as in the XMM-Newton pipeline, such candidates may never be submitted to Cash fitting in the first place.

The situation is slightly more complicated in the multi-band case. Because it is impossible to define a single source spectrum, it may be desirable to allow the flux in each band to be a free parameter – i.e. to fit the flux α_k in each k th band separately. Since one can easily conceive of a source which is bright in total but very faint in one of more spectral bands, it is entirely appropriate to retain negative values for $\hat{\alpha}_k$ s in the final source list. As mentioned above, all that this implies is that the flux of that source in that band is undetectable.

Clearly, a criterion to decide whether a given multi-band, free-spectrum detection should be retained or not should involve some function of the fitted fluxes in all bands. It is just not quite clear what this function should be. Simplest may be to simply add the fluxes and retain only those sources for which the result is positive.

The validity of the second objection to maximum likelihood fitting, namely that the background and signal amplitude are obtained from the same set of data, depends on the circumstances of the detection. However, in XMM-Newton practice at least, the fraction of image area occupied by identifiable sources is usually small. In this case it seems a reasonable procedure to make an estimate of background from those regions of the image which one has concluded probably don’t contain any detectable sources. The XMM-Newton pipeline performs a simple 2-step iterative procedure to calculate a background estimate by this method.

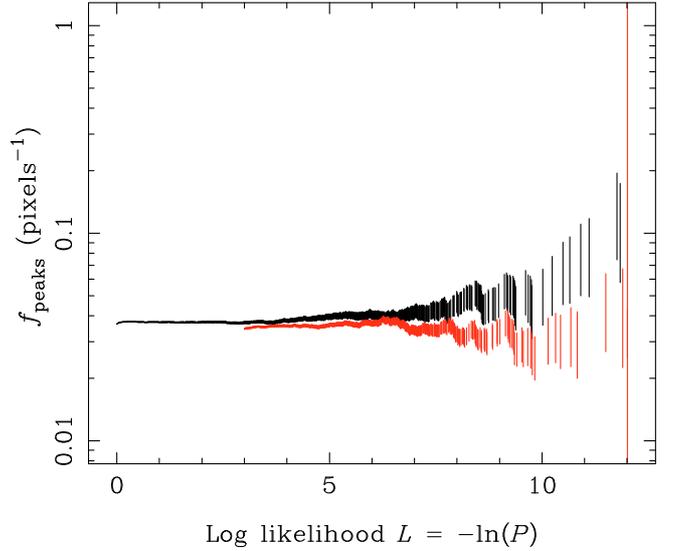


Fig. 4. A plot which allows the fraction of peak pixels f_{peaks} to be estimated. This value is just the height to which the curve asymptotes at small values of L . As described in the text, a flat slope for such a curve indicates that the shape (but not necessarily the normalization) of the empirical probability distribution agrees with that predicted by theory. The black bars show the results of the bespoke Cash fitting software; the half-tone bars (red in the electronic version) show the results of the SAS detection procedure. Both plots are truncated some way short of zero because only candidate sources having a “map” probability above a given cutoff were submitted to the fitting procedure. A number of computational reasons led to the choice of a much larger cutoff for the SAS fitting task (*emldetect*).

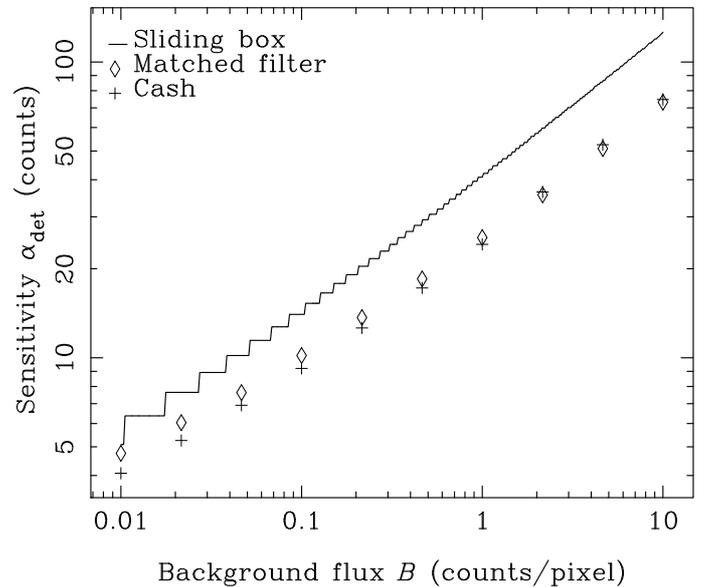


Fig. 5. Minimum detectable source counts α_{det} , plotted as a function of background flux B ; single-band case. The solid line gives the results of the sliding-box technique; the diamonds show the results of detection by matched linear filter; the crosses show the results of the Cash method.

Since one is in this case not using the same data for both background estimation and source fitting, the second objection does not apply.

Protassov et al. propose the replacement of the ML statistic in these problems by a Bayesian technique. While admittedly more rigorous, their technique would require at least a fresh

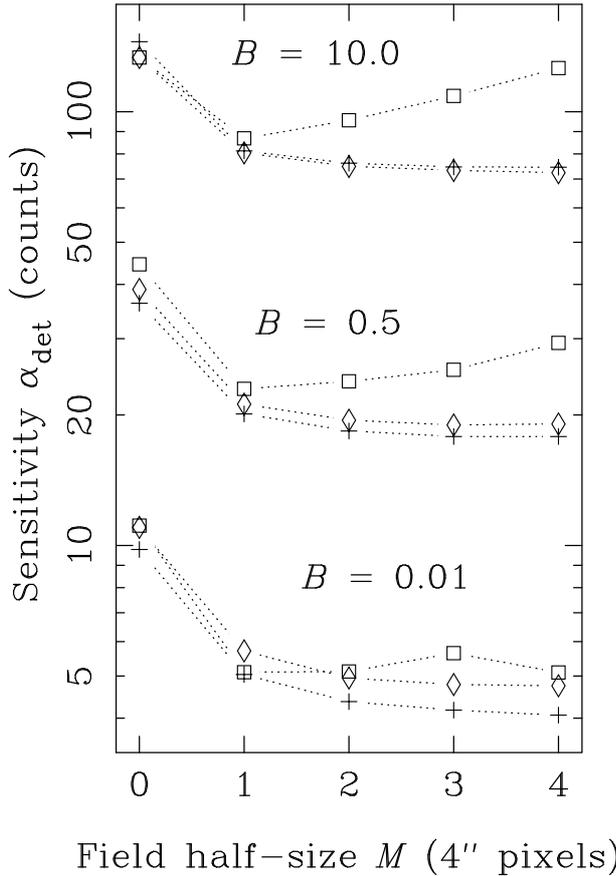


Fig. 6. The sensitivity of different detection methods as the size of the “detection field” is varied. The field side length is $2M+1$ pixels. Results are presented for three values of expected background flux B , given in counts per pixel. Those results belonging to the same method and the same value of B have been linked with dotted lines to improve the legibility. The squares represent the sliding box convolver; diamonds, the matched filter method; and crosses, the Cash method.

Monte Carlo to be performed for each source. Protassov et al. in their spectral line examples deal in significances on the order of a few percent. To calibrate the distribution function of their detection statistic, a Monte Carlo of a few hundred elements suffices. Source detection however typically demands a much more stringent significance level: 3×10^{-4} for 1XMM (Watson et al. 2003) for example. Monte Carlo ensembles with on order 10^5 or more members are required to calibrate distributions down to this level. Such a technique is not practical for the compilation of a source catalog with upward of 50 000 members.

2.4.4. Empirical verification of ML source detection

An empirical look at some Monte Carlo data shows indeed that the raw distribution of values of the Cash statistic does not appear to conform very closely to the ideal χ^2 curve. This is so whether one sets the value of U_{Cash} to zero whenever the fitted amplitude is negative (as in the examples of Protassov et al.) or whether one simply retains such values unaltered. However, discarding negative-amplitude fits from the pool of source candidates, as described in the preceding section, brings the distribution of the remainder very close to the χ^2 theoretical curve. Figure 3 shows the results of such a Monte Carlo.

For Figs. 3 and 4, it was not possible to deal only with a small patch of simulated-data pixels, such as for Figs. 1 and 2,

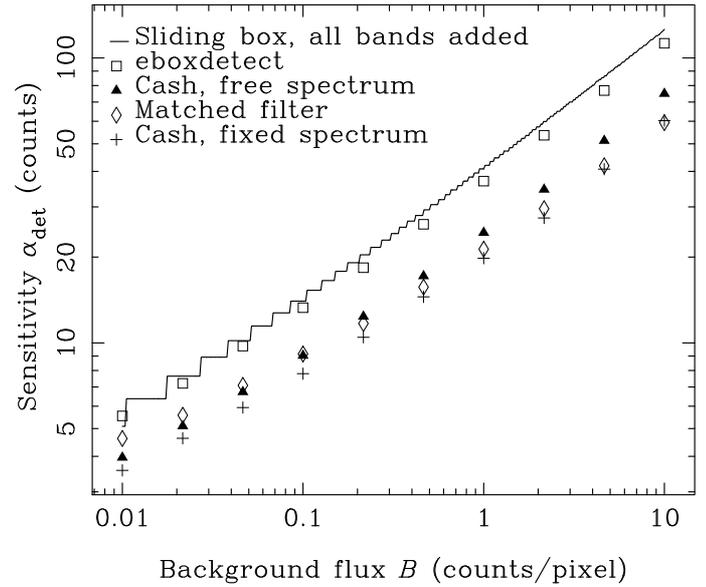


Fig. 7. Minimum detectable source counts α_{det} , plotted as a function of background flux B ; N -band case. The solid line gives the results of the technique in which all bands are summed before sliding-box is applied (method 1 this section). Squares represent the results of the *eboxdetect* sliding-box technique (method 2). Diamonds show the matched-filter results (method 3). The Cash method in which a source spectrum is assumed (method 4) is represented by crosses; the alternate, assumption-free Cash method (method 5) is represented by filled triangles.

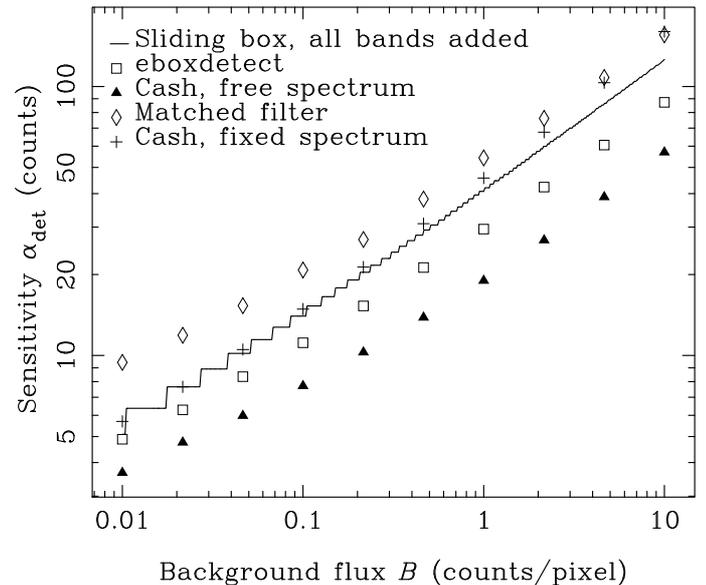


Fig. 8. Same as Fig. 7 except a hard source spectrum has been used to generate the data.

because in the present case it was necessary to perform all three steps of the generalized source-detection algorithm (Sect. 2.2). For this reason a much larger, circular field was used, with a radius (204 pixels) similar to that (~ 210 pixels) of the field of view of the XMM-Newton EPIC cameras, when expressed in the standard 4-arcsec pixels of the data products. No sources at all were used in the input model for the Monte Carlos of Figs. 3 and 4: all detections depicted in these figures are “false positives”.

The penalty for improving the shape of the distribution by discarding negative fits is that one needs to estimate the expected

fraction of such fits in order to normalize the distribution. When considering the statistics of “map”-stage Cash (i.e. fitting of amplitude alone at each pixel of an image), this fraction asymptotes to 0.5 in the limit of high background (see e.g. the zero intercept of the solid curve in Fig. 3). This value has been found to be an acceptable approximation throughout the range of background values used in the present paper.

Where the source position as well as amplitude is fitted using the Cash statistic, this simple approximation for the fraction of image pixels which yield candidates no longer applies. In this case one also needs to know the fraction of pixels which are peaks. There appears to be no way to estimate this apart from a Monte Carlo. One can derive the number a posteriori via a plot such as Fig. 4.

Figure 4 was created as follows. An ensemble of (false) source detections was created via a Monte Carlo as described above in the description of Fig. 3. The detection software makes an estimate of the probability P that a source of that amplitude or greater will occur by chance from background. If that estimate is proportional to the true probability, then the proportionality may be expressed as follows:

$$N(>P) \sim M_{\text{fields}} A_{\text{field}} f_{\text{peaks}} P$$

where $N(>P)$ is the number of sources in the ensemble having P or greater, M_{fields} is the number of fields, A_{field} is the area of the field in pixels, and f_{peaks} is the fraction of image pixels which are identified as peaks (i.e., positive-valued local extrema). (Note that the diamonds, circles and dotted line in Fig. 3 should all trend to f_{peaks} as U approaches zero.) Taking logs and manipulating gives

$$\log_{10} \left(\frac{N}{M_{\text{fields}} A_{\text{field}}} \right) + \frac{L}{\ln(10)} = \log_{10}(f_{\text{peaks}})$$

where the replacement $L = -\ln(P)$ has been made since this is the value written by *emldetect* to the column named DET_ML in the output source list.

For each point in Fig. 4, N is calculated as the number of sources having that value of L or higher. The error bars are just \sqrt{N} . Thus N increments as one proceeds from the highest L value to the lowest.

Results for two source-detection procedures are plotted in Fig. 4. The black bars show results from software written for the present paper which performed all three stages of the source detection (Sect. 2.2). The half-tone bars (red in the electronic version) result from use of the XMM-Newton SAS tasks *eboxdetect* followed by *emldetect*. Actual XMM-Newton product file headers were used in the simulated image, exposure map and background map files in order to facilitate acceptance of these by the SAS source detection tasks.

The asymptotic flatness of the Fig. 4 plots again supports the close correspondence between the actual distribution of values of the Cash statistic and the χ^2 distribution predicted by the Cash theory.

Finally, although it is beyond the scope of the present paper to perform detailed tests on *emldetect*, one can at least say from Fig. 4 that the value of DET_ML is being calculated correctly. This seems not to be consistent with the results shown in Fig. 4 of Brunner et al. (2008), which were also calculated from simulated data. However, their simulation was much more sophisticated: every field contained very many faint confusing sources, and a statistical procedure had to be used to identify “false positives”. This may explain the discrepancy with present results.

3. Sensitivity comparison

3.1. The model

The flux model used in the simulation was based on Eq. (2). The physical coordinates here of interest are the two spatial coordinates x and y , representing position on the detector (or, equivalently, position on a projection plane tangent to the celestial sphere), plus the energy E of the X-ray photons. Samples along both x and y axes are taken to be equally spaced, located indeed on an orthonormal array, centred on the source centre (x_0, y_0) ; but samples may be unevenly spaced in energy. It is convenient to modify Eq. (2) to make use of three indices, one per coordinate, viz:

$$e_{i,j,k} = B_{i,j,k} + \alpha S_{i,j,k}.$$

Both i and j are restricted to the range $-M$ to M for some small integer M . The symbol N is henceforth used to denote the number of energy bands.

The signal samples are related to the PSF $S(x, y, E)$ as follows:

$$S_{i,j,k} = \int_{x_i-\Delta x/2}^{x_i+\Delta x/2} \int_{y_j-\Delta y/2}^{y_j+\Delta y/2} dx dy S(x, y, E_k)$$

where Δx and Δy are the respective sizes of the pixels in the x and y directions and $E_k = (E_{k,\text{lo}} + E_{k,\text{hi}})/2$, i.e. the average of the bounding energies of band k . S is normalized such that

$$\frac{1}{\Delta x \Delta y} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} dx dy \sum_{k=1}^N S(x, y, E_k) = 1. \quad (4)$$

(Note that this is not the same normalization condition as used for Paper I.) Hence, α has units of X-ray counts and may be interpreted directly as the expected total number of counts, summed over all N bands and over the entire detector plane, which will be generated by the source S during a given exposure duration.

As for Paper I, both Δx and Δy are set to the same value of 4 arcsec, which is the same as the pixel size of the X-ray images produced as standard XMM-Newton products, and as used for source detection in both the 1XMM and 2XMM catalogs (Watson et al. 2003 and 2008).

The PSF used, for both single-band and multi-band calculations, was that corresponding to the pn camera of XMM-Newton on the optical axis, and at a photon energy of 1.25 keV, which is the centre of band 2 of the 1XMM catalog X-ray data products (Watson et al. 2003). In Paper I the PSF was allowed to vary in shape with energy band, as it is observed to do in XMM-Newton. This is discarded in the present paper in order to avoid confusion between the small effect due to the change in shape of S with energy and the large effect due to its changing magnitude.

The expected background flux $B_{i,j,k}$ in counts per pixel is taken throughout to be spatially invariant, but in general different within each energy band. Source and background spectra used, plus the associated energy band definitions, are taken from Table 1 of Paper I.

3.2. The sensitivity calculation

As noted in Sect. 2.3, the first step in source detection, thus also in calculating the sensitivity of that detection, is to decide the acceptable probability of a false detection. Here, as in Paper I, a value of $\exp(-8.0)$ is used for this probability cutoff P_{det} . This was the value used in the making of the 1XMM catalog (Watson et al. 2003).

The next step is to invert the appropriate expression for $P(>U; 0)$ to obtain U_{det} . In a minority of cases, an accurate closed-form expression for P is known: in these cases, P was inverted by use of Ridders' method (Ridders 1979) as described in Press et al. (2003). Ridders' method was chosen rather than a Newton-type method because of a paucity of expressions for $\partial P/\partial U$.

In the majority of cases, no closed-form expression or approximation for P is known which is accurate enough for present purposes. For each detection method without an accurate formula for P , a Monte Carlo technique was used to estimate U_{det} . An ensemble of η sets of $c_{i,j,k}$ were generated from a model comprising purely the background B_k . U was calculated for each set of counts values. The final ensemble of U_m values, for $m \in [1, \eta]$, was sorted into increasing order. U_{det} is then approximately that value of U_m for which $(\eta-m)/\eta \sim P_{\text{det}}$. Clearly η must be chosen such that $m \gg 1$.

Having obtained U_{det} , the final step is to invert an expression for $\langle U \rangle$ to derive α_{det} . An accurate and invertible expression for $\langle U \rangle$ could be found for the majority of the detection methods. Where not, a Monte Carlo technique had to be used for this step of the calculation as well. In this onerous and somewhat problematic procedure, a Ridders iteration was employed to converge on a value of α_{det} . For each trial value in the iteration sequence, an ensemble of sets of $c_{i,j,k}$ was obtained. Each set yielded its value of U . The mean of the ensemble of U values was tested against U_{det} and the trial α adjusted accordingly until convergence was judged to have been reached.

3.3. 1-band results

Three single-band detection methods are compared here, viz:

1. Sliding-box: U is obtained simply by summing the c_i values within the box.
2. Optimized linear filter: U is obtained via a weighted sum of the c_i , the weights being chosen to optimize the sensitivity.
3. Cash-statistic detection: U given by the Cash statistic where amplitude only is fitted.

Formulas for $P(>U; 0)$ and $\langle U \rangle$ are given in Appendix B.1.

Sensitivities of the three methods are compared for a range of values of background flux B in Fig. 5. All values were calculated for a square array of pixels with $M = 4$.

The reason the sliding-box curve shows stepwise discontinuities is because (as described in Sect. 2.3) U for any method of detecting signals in Poissonian data must be a discrete function, not defined for all values on the real line, due to the discrete nature of the counts value in any pixel. All the P functions, and any expression derived from them, thus descend in step fashion. This is most obvious for the sliding-box method because U for this method is highly degenerate – i.e. there are many arrangements of counts c_i which will produce the same U . The steps are still present for the other methods but are usually too finely spaced to be noticeable.

As one might expect, both the matched-filter and Cash methods are seen to be more sensitive (by a factor of at least 20% at any value of background) than the sliding-box method. This is because they both make use of more information, namely the shape of S , to better discriminate between S and B . What is not so expected is that the Cash method appears to be marginally better than the matched-filter method. At first it seemed possible that this was an artifact, a result of insufficient convergence of the Powell's-method routine which calculates the weights for the matched-filter method; however, tests using much stricter convergence criteria failed to reveal any significant dependence. The

matched filter detection is by definition the optimum *linear* filter; however the Cash detection algorithm is based however on a *nonlinear* expression, which is optimized not just for the general conditions but for the particular arrangement of counts within each $(2M + 1)$ by $(2M + 1)$ detection field. One must assume that this accounts for its better performance.

Figure 5 is comparable to Fig. 6 of Paper I, with the addition of the points for the Cash results. Some differences can however be observed between the sensitivity values displayed in Fig. 5 and those given in the earlier figure. There are three reasons for this. Firstly, the normalization expression for S is different: in Paper I, the normalization integral for S extended only as far as the boundary of the square $(2M + 1)$ by $(2M + 1)$ patch of pixels in which counts were considered; in the present paper the integral extends over the whole plane.

Secondly, a smooth approximation to $P(>U_{\text{box}}; 0)$ was used in Paper I, rather than the true step-function form used here.

Thirdly, whereas in the present paper the same value of $M = 4$ has been used for all methods, in Paper I the values for the sliding-box method were calculated using $M = 2$. Figure 4 in Paper I was provided to allow the reader to estimate the effect of this choice. This figure is repeated here as Fig. 6, with the Cash results included.

The points to note from Fig. 6 are twofold. The first is that, for both the matched-filter and Cash methods, the minimum detectable source count always decreases as the size of the field is increased, whereas the sliding box method reaches a minimum, with a slight dependence on the background level. The reasons for this were explored in Paper I but, to reprise briefly: when confronted with a large field, most of which must be almost pure background, the sliding box method adds in all the background-dominated counts with equal weight, which therefore increases the total noise; whereas the matched-filter method assigns smaller weights to pixels as the fraction of S to B decreases. The same is true of the Cash method, for which the fit is dominated by pixels having a high S/B ratio.

The second point to observe in Fig. 6 is that all methods become of roughly similar worth as M becomes small. For $M = 0$, corresponding to a detection field consisting of a single pixel, one would expect them to become identical, since any knowledge of the shape of S is useless in this case. The reason for the persistence of some differences at $M = 0$, particularly at low background, is unknown. However, the differences occur within a regime for which the expected total number of counts is less than 1. One might expect most measures to be a little “noisy” under such circumstances.

3.4. N-band results

Five multi-band detection methods are compared here, viz:

1. Sliding-box: U is obtained simply by summing the c_i values within the box, over all bands.
2. *eboxdetect*: this is also a sliding box method, but differs from the previous in that calculation of U is two-stage. The first stage calculates U_k for each k th band by summing values as usual. U is then obtained from the U_k via a formula (see Appendix B.2.2) for which the theoretical underpinning is unclear.
3. Optimized linear filter: U is obtained via a weighted sum of the c_i over all pixels and bands, the weights being chosen to optimize the sensitivity. The complete form of S , thus of the source spectrum, must be assumed.

4. Cash detection, fixed spectrum: U is given by the Cash statistic where amplitude is fitted for all bands in parallel. A source spectrum must be assumed. There is thus only 1 degree of freedom in the fit.
5. Cash detection, free spectrum: same as the preceding, except that the amplitude is fitted for each band separately, giving N degrees of freedom, for N bands.

Formulas for $P(>U; 0)$ and $\langle U \rangle$ are given in Appendix B.2.

Sensitivities of the five N -band methods are compared for a range of values of background flux B in Fig. 7. All values were calculated for a square array of pixels with $M = 4$.

The background flux B in all figures relating to multi-band detection was calculated by summing the background B_i in each band. This provision, together with the normalization of S described in Eq. (4), allows one to compare directly single- to multi-band detection. This was not the case in Paper I.

The same changes and improvements as described in Sect. 3.3 hinder direct comparison of Fig. 7 of the present paper with Fig. 9 of Paper I. However if one increases all values in Fig. 9 Paper I by about 40% to compensate for the different normalization of S in that paper, the respective curves (diamonds in both papers indicate the matched-filter results, whereas the Monte Carlo-corrected *eboxdetect* results are indicated by crosses in Paper I and squares in the present paper) are seen to match reasonably well.

The normalization condition for S chosen in the present paper does however facilitate comparison between the 1-band and N -band situations. The summed, sliding-box detection method is seen, as expected, to provide identical results here as in the single-band case. Some of the other N -band methods are conceptually similar to 1-band counterparts, and have been given the same graph symbols to highlight this fact; however, only for the sliding-box method are the N -band and 1-band sensitivity values expected to be identical.

The first thing to note is that the *eboxdetect* version of sliding-box, in which the likelihoods for each band are summed to give U , performs slightly better than the simplistic summed-band version. Why? The answer is that the simple sum actually implies an assumption about the source spectrum, namely that it is flat – or, more precisely, that the source spectrum is such that the optimum weights per band are identical. This method might therefore be expected to perform better than the *eboxdetect* method on sources which obey this criterion, and worse (as in the present case) when they do not. However there seemed little point to the present author in testing this speculation because the summed-likelihood sliding box method is in any case far from optimum.

Again the equivalent matched-filter and Cash methods, namely methods 3 and 4 above, represented respectively by diamonds and crosses, perform about the same. The improved performance of the Cash method at low values of background is no doubt due to the same reasons as in the single-band case.

As one might expect, the version of the Cash statistic which does not require an a priori decision about the source spectrum performs less well than methods which do, at least in the present case in which the true source spectrum matches the template. The contrary case is examined in the next subsection.

3.5. Effect of non-matching spectra on sensitivity

Figure 8 is similar to Fig. 7, except that there is now a mismatch between the spectrum used to generate the counts (and

to calculate the sensitivity) and the spectrum assumed in the detection process. Such a mismatch is of course only meaningful for those methods which require a spectrum to be assumed. The hard source profile in Table 1 of Paper I was used to generate the Monte Carlo data, whereas the more representative source spectrum, the same used for both purposes in Fig. 7, has been used in the appropriate detection schemes.

The summed-band sliding-box method again provides the unvarying benchmark. Among the remaining methods, two trends are evident: those methods (respectively *eboxdetect* and free-spectrum Cash) which do not assume a source spectrum perform about the same as previously; whereas those which do (fixed-spectrum Cash and matched-filter) perform now significantly worse. This is as one would expect.

It is interesting that the fixed-spectrum Cash and matched-filter methods now diverge quite markedly at low values of background. The reason for this behaviour is not known.

4. Some $\langle U \rangle$ and P difficulties

4.1. An approximate expression for $\langle U_{\text{Cash}} \rangle$

As described in Appendices B.1.3, B.2.4 and B.2.5, no closed-form expression for $\langle U_{\text{Cash}} \rangle$ is known. In its absence, inversion to derive α_{det} from U_{det} must proceed via a tedious and problematic Monte Carlo. Such an approach is clearly going to be inadequate for “in the field” calculations of sensitivity. The following expression appears however to offer an acceptable approximation for such cases:

$$\langle U_{\text{Cash}} \rangle \sim 1 + 2 \sum_i \left([B_i + \alpha S_i] \ln \left[\frac{B_i + \alpha S_i}{B_i} \right] \right) - 2\alpha \sum_i S_i. \quad (5)$$

This expression is compared to Monte Carlo data for several values of background in Fig. 9.

Equation (5) was obtained by a series of guesses at reasonable approximations. Initially, a formula was sought for the case in which negative-amplitude fits are *not* discarded. The additional “fudge factor” of unity produced a function which gave an excellent fit to the Monte Carlo data for such a case. Discarding negative fits, as is done for all results in the present paper, degrades this close match at low amplitudes. However, the approximation is seen to be a good match at α values corresponding to realistic sensitivities (the half-tone diamonds in Fig. 9). For this reason it is considered to be useful for the construction of sensitivity maps for Cash source detection.

4.2. A Cash statistic for sliding-box

There are two problems in the calculation of $\langle U \rangle$ and P for the *eboxdetect* multi-band algorithm: firstly, there appears to be no closed-form expression for $\langle U \rangle$; secondly, the best available approximation for P is significantly inaccurate at low values of background.

4.2.1. Incorrect $\langle U \rangle$

In Paper I, sensitivities for the *eboxdetect* detection method were calculated using the assumption that $\langle U \rangle$ could be approximated by

$$\langle U \rangle \sim -2 \sum_{k=1}^N \ln(P[>U_k; \alpha=0]). \quad (6)$$

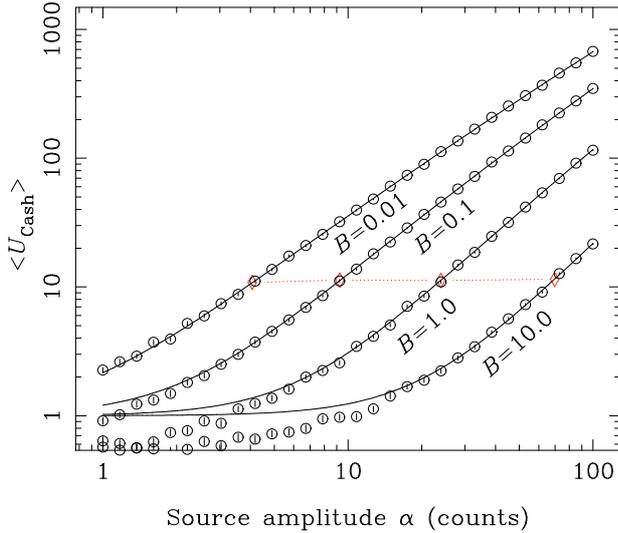


Fig. 9. Comparison between the approximate expression for $\langle U_{\text{Cash}} \rangle$ developed in Sect. 4.1 and Monte Carlo estimations of it. The Monte Carlo points are indicated by circles centred on error bars (many of the error bars are too small to be visible on this plot). All values were calculated for the single band case with $M = 4$. Half-tone diamonds, connected by dotted lines (both coloured red in the electronic version), indicate the sensitivities appropriate to each background value.

$P(>\langle U_k \rangle; \alpha=0)$ here is given by

$$P(>\langle U_k \rangle; \alpha=0) = 1 - Q\left(\sum_{i,j} B_{i,j,k} + \alpha \sum_{i,j} S_{i,j,k}, \sum_{i,j} B_{i,j,k}\right)$$

where Q is the complementary incomplete gamma function¹ defined by

$$Q(a, x) = \frac{1}{\Gamma(a)} \int_x^\infty dt \exp(-t) t^{a-1}.$$

This expression was numerically inverted by use of Ridder's method to derive the sensitivity α_{det} from U_{det} . Monte Carlo experiments show, however, that Eq. (6) overestimates α_{det} by about 10% over most of the background range. It is simply incorrect and should not be used.

4.2.2. Inaccurate P

Figure 3 of Paper I shows that use of the *eboxdetect* formula for P (i.e., the formula used to calculate the values written by *eboxdetect* to the LIKE column of its output source list) can be incorrect by as much as an order of magnitude at moderately low values of background. Without an understanding of the theoretical underpinnings of the *eboxdetect* method, it is difficult to suggest a rigorously better alternative. All sensitivity values calculated for this method in the present paper therefore made use of a Monte Carlo approximation.

¹ The somewhat awkward notation $1 - Q$ is here chosen, instead of the more compact and natural symbol $P = 1 - Q$, representing the other incomplete gamma function, merely in order to avoid using P for more than one quantity. Other alternative representations for the incomplete gamma functions alas offer no improvement since they involve additional mentions of the gamma function $\Gamma(a)$.

4.2.3. A more accurate alternative

One way around these difficulties is to recast the *eboxdetect* algorithm in terms of the Cash statistic. This probably doesn't make the method any more sensitive, but it does make its sensitivity (and its detection likelihoods) easier to calculate.

Since we are not using the information in the spatial variation of S , both background and counts can be added across the detection field to give, respectively, a single number for each band. We define U_{BoxCash} as

$$U_{\text{BoxCash}} = \sum_{k=1}^N U_k$$

where U_k has the usual Cash form, even though there is here only one "pixel" to sum over:

$$U_k = 2 c_{\text{total},k} \ln\left(\frac{B_{\text{total},k} + \hat{\alpha}_k S_{\text{total},k}}{B_{\text{total},k}}\right) - 2\hat{\alpha}_k S_{\text{total},k}$$

for

$$B_{\text{total},k} = \sum_{i=-M}^M \sum_{j=-M}^M B_{i,j,k}$$

etc. In this instance it is easy to deduce that, for the k th band,

$$\hat{\alpha}_k = \frac{c_{\text{total},k} - B_{\text{total},k}}{S_{\text{total},k}};$$

The expression for U_k thus becomes

$$U_k = B_{\text{total},k} - C_{\text{total},k} \ln\left(\frac{C_{\text{total},k}}{B_{\text{total},k}}\right) \text{ where } C_{\text{total},k} > 0 \\ = B_{\text{total},k} \text{ otherwise.}$$

P is as given in the Appendix B.2.5; and we can now use the approximation in Eq. (5) to estimate (and therefore invert) $\langle U \rangle$. Monte Carlos show that the Cash χ^2 prediction remains a good fit to P , with the usual $\times 0.5$ normalization.

5. Conclusions

This paper has attempted to set out clearly the issues involved in the detection of sources in images which exhibit Poisson noise. A three-stage detection algorithm was proposed in which the first stage makes a map of the detection likelihood at each image pixel, the second searches for peaks in this map and the third fits a source profile to each peak so as to return best-fit values and uncertainties for the amplitude and position of each source. Some conditions which the data must obey in order for this approach to succeed were examined. The proper way to define detection sensitivity was also rather carefully gone into.

It was realized after the issue of Paper I that the probability distribution of map values is not in general the same as the distribution of values at a set of pixels restricted to peaks in that map. This has implications for the calculation of the sensitivity of a detection method, since the starting point in this calculation, namely the acceptable maximum in the number density of false detections returned by the method, translates to a cutoff value in the peak rather than the map distribution. On the other hand, because the map and peak distributions can be shown to converge at high values, it was felt allowable to use the map distribution as a proxy for the peak one when comparing sensitivities

of different detection methods. Thus the results of Paper I are not invalidated and could be extended here.

The Cash likelihood-ratio statistic (LRS) as a means of detecting sources amid Poissonian noise was here described. This statistic has been used for source detection at least in Rosat images (Cruddace et al. 1977; and Boese & Doebereiner 2001) and also for XMM-Newton (Watson et al. 2003, 2008). According to Cash (1979), or rather to Wilks (1963) as cited by Cash, such an LRS ought to be distributed as χ^2 with q degrees of freedom, where q is the number of free parameters. It has however been claimed that the LRS is not well suited to source detection because the probability distribution of values of such a statistic calculated under the null hypothesis (i.e. that the image contains no sources) is not well known. Protassov et al. (2002) for example maintain that the source-detection application of LRS disobeys some of the necessary regularity conditions under which the distribution of an LRS can be expected to follow that of χ^2 . In the present paper this claim is disputed; Monte Carlo experiments are also described which show, under typical X-ray instrument conditions, and provided that sources with negative fitted amplitude are discarded from the ensemble, that the distribution of the Cash statistic, evaluated for images simulated from a source-free flux model, adheres in shape at least quite closely to a χ^2 distribution with the appropriate degrees of freedom. All that is necessary is to evaluate the normalization constant between the two.

Although it was found convenient for purposes of the present paper to write bespoke code to perform the Cash-statistic source detection and fitting, the XMM-Newton SAS task *emldetect* was briefly compared to this bespoke code, and its results shown to also obey the χ^2 distribution.

The main thrust of the present paper has been to compare the sensitivity of different detection methods. The Cash-statistic detection is shown to be the best of the methods tested. In terms purely of sensitivity it is seen to be only marginally superior to an optimized linear filter; however the Cash method is shown to possess two further advantages. Firstly, it is more flexible, since it allows the user the option of assuming a source spectrum or leaving this free. Secondly, for the final detection probability, one need not just accept the “map” value at the pixel nearest to the final source position: instead one can calculate a final, accurate value specific to that fitted position.

The flexibility of the Cash statistic was also demonstrated by showing how it can be adapted to calculate detection likelihoods for sliding-box source detection as performed for example by the SAS task *eboxdetect*. The advantage in doing so is that one can make use of the demonstrated close agreement between the distribution of the Cash statistic and χ^2 . This would represent a considerable improvement over the likelihood formula employed within *eboxdetect*.

As a final result, the paper describes approximations which can be used to calculate sensitivity maps appropriate to Cash-style source detection. For XMM-Newton the only previous ways to do this have been via the SAS task *esensmap* (which examination of the code shows contains an algorithm which could only have serendipitous success), or via an empirical correction to sliding-box sensitivities (e.g. Carrera et al. 2007).

Acknowledgements. Thanks are due firstly to Georg Lamer for providing the vital clue to understanding the Cash algorithm; and secondly to Anja Schröder for valuable assistance in obtaining XMM-Newton data.

References

- Bergmann, U. C., & Riisager, K. 2002, *Nuc. Inst. Meth. Phys. Res. A*, 489, 444
 Boese, F. G., & Doebereiner, S. 2001, *A&A*, 370, 649
 Brandt, W. N., & Hasinger, G. 2005, *ARA&A*, 43, 827
 Brunner, H., Cappelluti, N., Hasinger, G., et al. 2008, *A&A*, 479, 283
 Carrera, F. J., Ebrero, J., Mateos, S., et al. 2007, *A&A*, 469, 27
 Carter, J. A., & Read, A. M. 2007, *A&A*, 464, 1155
 Cash, W. 1979, *ApJ*, 228, 939
 Cruddace, R. G., Hasinger, G. R., & Schmitt, J. H. M. M. 1977, *Astronomy from Large Databases*, Garching, 177
 DePonte, J., & Primini, F. A. 1993, *Astronomical Data Analysis Software and Systems II*, ASP Conf. Ser., 52, 425
 Dobrzycki, A., Jessop, H., Calderwood, T. J., & Harris, D. E. 2000, *BAAS*, 32, 1228
 Freeman, P. E., Graziani, C., Lamb, D. Q., et al. 1999, *ApJ*, 524, 753
 Gilli, R., Comastri, A., & Hasinger, G. 2007, *A&A*, 463, 79
 Koyama, K., Tsunemi, H., Dotani, T., et al. 2007, *PASJ*, 59, S23
 Kuntz, K. D., & Snowden, S. L. 2008, *A&A*, 478, 575
 Lyons, L. 2007, *ASP Conf. Ser.*, 371, 361
 Mattox, J. R., Bertsch, D. L., Chiang, J., et al. 1996, *ApJ*, 461, 396
 Nakamoto, M., Minamihara, H., & Ohta, M. 2000, *Elec. Comm. Japan*, 83, 20
 Pilla, R., Loader, C., & Taylor, C. C. 2005, *Phys. Rev. Lett.*, 95, 230202
 Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 2003, *Numerical Recipes in Fortran 77: the Art of Scientific Computing*, 2nd Ed. (Cambridge UK: Cambridge University Press), 1
 Protassov, R., van Dyk, D. A., Connors, A., Kashyap, V. L., & Siemiginowska, A. 2002, *ApJ*, 571, 545
 Ridders, C. J. F. 1979, *IEEE Transactions on Circuits and Systems*, CAS-26, 979
 Serlemitsos, P. J., Soong, Y., Chan, K.-W., et al. 2007, *PASJ*, 59, S9
 Stewart, I. M. 2006, *A&A*, 454, 997
 Strüder, L., Briel, U., Dennerl, K., et al. 2001, *A&A*, 365, L18
 Turner, M. J. L., Abbey, A., Arnaud, M., et al. 2001, *A&A*, 365, L27
 Watson, M. G., Pye, J. P., Denby, M., et al. 2003, *Astron. Nachr.*, 324, 89
 Watson, M. G., Schroeder, A. C., Fyfe, D., et al. 2008, *The XMM-Newton Serendipitous Survey*, VI. The Second XMM-Newton Serendipitous Source Catalogue, *A&A*, submitted
 Weiskopf, M. C., Brinkman, B., Canizares, C., et al. 2002, *PASP*, 114, 1
 Wilks, S. S. 1963, *Mathematical Statistics* (Princeton: Princeton University Press), Chap. 13

Appendix A: Cash amplitude fitting

To calculate a map of the Cash statistic (Sect. 2.4.1) in a single spectral band, we centre a PSF S on each pixel of the image and fit only its amplitude α . The fit is performed over a small number N pixels surrounding the pixel for which we are calculating the value. L_{null} is calculated by setting α to zero. The resulting probability is then the probability of obtaining the fitted value $\hat{\alpha}$ when there is in fact nothing there but background B . The expression, derived from Eq. (3), for the Cash statistic relevant to this case is

$$U_{\text{Cash}} = 2 \sum_{i=1}^N c_i \ln \left(\frac{B_i + \hat{\alpha} S_i}{B_i} \right) - 2\hat{\alpha} \sum S_i \quad (\text{A.1})$$

where $\hat{\alpha}$ is the fitted value of α .

As mentioned in Sect. 2.4.3, statistically speaking it is permitted for $\hat{\alpha}$ to be negative; whether we later discard such occurrences depends on the physical model which is appropriate, and also on what we wish to do with the data. What is *not* permitted is for $B_i + \hat{\alpha} S_i$ to be negative for any i . (If the recorded counts at the i th pixel is >0 , $B_i + \hat{\alpha} S_i$ is not permitted to equal zero either – but this is a formal condition which has no effect on the computational practicalities.) The absolute limit on permitted $\hat{\alpha}$ values in the negative direction is thus $-\min(B/S)$.

If $c_i = 0$ for all N pixels in the fitted patch, $\hat{\alpha}$ should be set to this limit. This, the model with the least amount of flux which the parameter limits permit, is the one most likely to explain the lack of counts. If some $c_i > 0$ though, of course $\hat{\alpha}$ may be larger

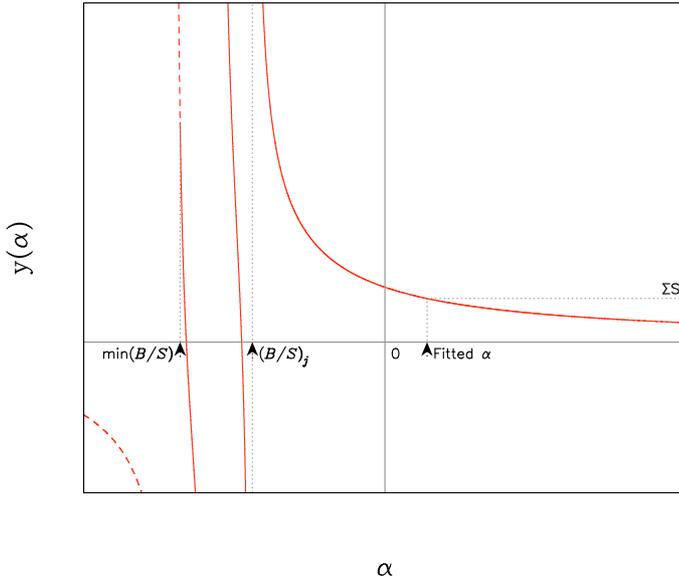


Fig. A.1. An example plot of Eq. (A) with 3 singularities. The dashed portion represents values beyond the “statistical limit”.

than $-\min(B/S)$. This possibility can be explored by taking the derivative of U with respect to $\hat{\alpha}$ and setting it to zero. This gives

$$\frac{\partial U_{\text{Cash}}}{\partial \hat{\alpha}} = 2 \sum_{i=1}^N \left(\frac{c_i}{\hat{\alpha} + B_i/S_i} \right) - 2 \sum S_i = 0. \quad (\text{A.2})$$

Equation (A.2) can be solved numerically for $\hat{\alpha}$, using for example the fast and robust Newton-Raphson method (cf. Press et al. 2003). To help in visualizing how to apply this it is useful to look at a diagram. Figure A.1 shows an example plot of the function

$$y(\hat{\alpha}) = \sum_{i=1}^N \left(\frac{c_i}{\hat{\alpha} + B_i/S_i} \right).$$

y has M hyperbolic singularities, one for each of the subset of M values of c_i which are >0 . Although formally speaking there are M solutions to Eq. (A.2) (indicated in Fig. A.1 by intersections of the graph with the horizontal line at $y = \sum S_i$), the largest solution is clearly the one desired.

Newton-Raphson solution can be facilitated in two ways: by estimating bounds for $\hat{\alpha}$ and by performing a coordinate transform. Firstly, let us label as the j th that pixel for which the following conditions hold: firstly that $c > 0$ for that pixel and secondly that, of that subset of M pixels for which this holds, B/S is the smallest. Clearly the most positive singularity occurs at $\hat{\alpha} = -(B/S)_j$. $\hat{\alpha}$ is bounded as follows:

$$\frac{c_j}{\sum S_i} \leq \hat{\alpha} + B_j/S_j \leq \frac{\sum c_i}{\sum S_i}.$$

The rate of convergence of the Newton-Raphson algorithm can be accelerated by performing a coordinate transform

$$u = 1/(\hat{\alpha} + B_j/S_j).$$

The “statistical” bounding condition on $\hat{\alpha}$, namely $\hat{\alpha} > -\min(B/S)$, must be applied if the Newton-Raphson solution exceeds it, as is possible.

Appendix B: Formulae for sensitivity calculations

B.1. 1-Band detection

Expressions for U , $P(>U; 0)$ and $\langle U \rangle$ are given below for the three single-band detection methods compared in Sect. 3.3. The k index is superfluous in the present single-band scenario and has therefore been suppressed for the time being.

B.1.1. 1-band sliding-box

$$U = \sum_{i=-M}^M \sum_{j=-M}^M c_{i,j}. \quad (\text{B.1})$$

$$P(>U; 0) = 1 - Q(U, [2M + 1]^2 \times B_{i,j}),$$

where Q is the complementary incomplete gamma function as previously.

$$\langle U \rangle = \sum_{i=-M}^M \sum_{j=-M}^M B_{i,j} + \alpha \sum_{i=-M}^M \sum_{j=-M}^M S_{i,j}.$$

B.1.2. 1-band optimized linear filter

$$U = \sum_{i=-M}^M \sum_{j=-M}^M w_{i,j} c_{i,j},$$

where the $w_{i,j}$ are optimum weights, calculated by a procedure described in Paper I.

$$P(>U; 0) \sim Q\left(\frac{B'}{w_{\text{equiv}}}, \frac{U}{w_{\text{equiv}}}\right)$$

where

$$B' = \sum_{i=-M}^M \sum_{j=-M}^M w_{i,j} B_{i,j}$$

and

$$w_{\text{equiv}} = \frac{\sum_{i=-M}^M \sum_{j=-M}^M w_{i,j}^2 B_{i,j}}{\sum_{i=-M}^M \sum_{j=-M}^M w_{i,j} B_{i,j}}$$

$$\langle U \rangle = \sum_{i=-M}^M \sum_{j=-M}^M w_{i,j} (B_{i,j} + \alpha S_{i,j}).$$

B.1.3. 1-band Cash-statistic detection

Expansion of Eq. (A.1) to include the extra spatial index gives

$$U = 2 \sum_{i=-M}^M \sum_{j=-M}^M c_{i,j} \ln\left(\frac{B_{i,j} + \hat{\alpha} S_{i,j}}{B_{i,j}}\right) - 2\hat{\alpha} \sum_{i=-M}^M \sum_{j=-M}^M S_{i,j}. \quad (\text{B.2})$$

$$P(>U; 0) \sim Q(1/2, U/2).$$

There appears to be no easy route to an exact, closed-form expression giving $\langle U_{\text{Cash}} \rangle$ in terms of B , α and S (although see Sect. 4.1 for an approximate formula). One can estimate the α appropriate to a given value of $\langle U \rangle$ via iterative Monte Carlos as described in Sect. 3.2; in fact this must be done if one wants to investigate the effect on the Cash detection algorithm of fitting with the wrong shape S . All the sensitivity values for the Cash algorithm presented in the present paper have therefore been estimated by this method.

B.2. N -band detection

Expressions for U , $P(>U; 0)$ and $\langle U \rangle$ are given below for the five single-band detection methods compared in Sect. 3.4.

B.2.1. Sliding-box detection on a single, all-band image

This is exactly the same as the algorithm described in Sect. B.1.1. The sum over N bands is here just made explicit.

This method does not require the user to choose a source spectrum template.

$$U = \sum_{i=-M}^M \sum_{j=-M}^M \sum_{k=1}^N c_{i,j,k},$$

$$P(>U; 0) = 1 - Q\left(U, [2M+1]^2 \times \sum_{k=1}^N B_{i,j,k}\right),$$

and

$$\langle U \rangle = \sum_{i=-M}^M \sum_{j=-M}^M \sum_{k=1}^N B_{i,j,k} + \alpha \sum_{i=-M}^M \sum_{j=-M}^M \sum_{k=1}^N S_{i,j,k}.$$

B.2.2. Sliding-box detection, *eboxdetect* statistic

This is the detection method employed by the XMM-Newton SAS task *eboxdetect* (online description at <http://xmm.vilspa.esa.es/sas/current/doc/eboxdetect.ps.gz>). The method does not require the user to choose a source spectrum template.

U is calculated in a three-step process. In the first step, counts within the box are summed for each k th band to give U_k . The second step forms $P(>U_k; 0)$ for each k . The final U is then given by

$$U = \sum_{k=1}^N L_k \quad (\text{B.3})$$

where

$$L_k = -2 \ln(P[>U_k; \alpha = 0]).$$

No description of *eboxdetect* appears to have been published; hence the theoretical basis of this expression is obscure.

If one assumes (which is not strictly true) that the L_k above follow distributions which don't change with k , one can apply the Central Limit Theorem to Eq. (B.3) to give, in the limit of high N , the following form² for P :

$$P(>U; 0) = 0.5 \left(1 + Q[0.5, U^2/2N\sigma_L^2]\right)$$

where σ_L^2 is the variance of L_k for any k . However in *eboxdetect* practice, still evident in the code in *eboxdetect* version 4.19, the actual expression used is

$$P(>U; 0) \sim Q(N, U/2). \quad (\text{B.4})$$

Some experimentation in both Paper I and the present paper confirms that Eq. (B.4) does appear to match Monte Carlo results

in the limit of high background. The issue is in any case academic for present purposes, since as in all cases where only an approximate expression for P is available, the results presented here are derived by use of a Monte Carlo estimation of P .

No closed-form expression for $\langle U \rangle$ exists: it must be calculated numerically.

B.2.3. N -band optimized linear filter

Essentially this is the same algorithm as described in Sect. B.1.2, just with the sum over the N bands made explicit. This method requires the user to choose a source spectrum template, in order to calculate the optimum weights w .

$$U = \sum_{i=-M}^M \sum_{j=-M}^M \sum_{k=1}^N w_{i,j,k} c_{i,j,k};$$

P is the same, but with of course

$$B' = \sum_{i=-M}^M \sum_{j=-M}^M \sum_{k=1}^N w_{i,j,k} B_{i,j,k}$$

and

$$w_{\text{equiv}} = \frac{\sum_{i=-M}^M \sum_{j=-M}^M \sum_{k=1}^N w_{i,j,k}^2 B_{i,j,k}}{\sum_{i=-M}^M \sum_{j=-M}^M \sum_{k=1}^N w_{i,j,k} B_{i,j,k}};$$

$$\langle U \rangle = \sum_{i=-M}^M \sum_{j=-M}^M \sum_{k=1}^N w_{i,j,k} (B_{i,j,k} + \alpha S_{i,j,k}).$$

B.2.4. Cash-statistic detection with a spectrum template

Essentially this is the same algorithm as described in Sect. B.1.3, just with the sum over the N bands made explicit. This algorithm requires the user to choose a source spectrum template.

$$U = 2 \sum_{i=-M}^M \sum_{j=-M}^M \sum_{k=1}^N c_{i,j,k} \ln \left(\frac{B_{i,j,k} + \hat{\alpha} S_{i,j,k}}{B_{i,j}} \right) - 2\hat{\alpha} \sum_{i=-M}^M \sum_{j=-M}^M S_{i,j}.$$

P is the same as in Sect. B.1.3, and $\langle U \rangle$ is again estimated via a Monte Carlo procedure (though see Sect. 4.1).

B.2.5. Cash-statistic detection without a spectrum template

The difference between this method and the preceding one is that here a value of $\hat{\alpha}$ is calculated, i.e. the source PSF is fitted to the data, separately for each band. No knowledge of the source spectrum is required. U is then formed by summing the U_k for each band. Since there are now N degrees of freedom in the fit, P is here given by

$$P(>U; 0) \sim Q(N/2, U/2).$$

$\langle U \rangle$ is again estimated via a Monte Carlo procedure (though see Sect. 4.1).

² This derivation makes use of the identity $\text{erfc}(x) = Q(0.5, x^2)$.