DR01/02 Topical Report: Background subtraction from imaging data

Revised version
8th October 2010

1 Introduction

This document reports on the development of a sky subtraction module for the reduction of imaging data from ESO's near-infrared instruments, HAWK-I and ISAAC, as well as optical instruments such as VIMOS and WFI. Some physical background on the characteristics of NIR sky emission is given in Sect. 2. The requirements for the software module are discussed in Sect. 3. Sect. 4 briefly lists available software packages and their approach to sky subtraction. In Sect. 5 and 6 we describe the algorithms chosen to be implemented in the ESOsoft module and discuss several approaches to aspects of the algorithm, in particular concerning the rejection of outliers and masking of objects. In Sect. 7 we discuss a number of methods for stacking images, essentially robust estimators for the average of a sample of pixel values. These methods are used for the computation of the background correction and will also be useful for the stacking of registered images later in the imaging data reduction workflow. We also discuss how noise in the input images is transferred to the output data. A brief description of the software package follows in Sect. 8. We finally present results from the application of the code to several test data sets from HAWK-I and ISAAC in Sect. 9.

2 Physical background

Even at the best sites for astronomical observations, the night sky is not really dark but contributes significantly to the light flux registered by astronomical imaging instruments. In the optical wave-bands, there is a very faint continuum, caused by scattered light from terrestrial light sources or the moon. In addition, there are isolated emission lines (in particular $[\mathrm{O\,I}]\,\lambda\,5577$ and $\mathrm{Na\,I}\,\lambda\lambda\,5890,5896$) in the blue bands (e.g. B, V) and the OH molecular bands in the redder bands (Osterbrock et al. 1996, and references therein). The overall surface brightness is fairly low (Table 1. In the blue bands (including R for most instruments), the spatial structure is smooth; in typical deep extragalactic observations the sky varies on scales that are larger than the size of the largest objects in the field. The strength of the night sky emission lines and the surface brightness of the night sky vary somewhat over the course of the night and from night to night.

In the red and infrared bands (e.g. I and Z and sometimes R), the interference of the emission lines in the thin layers of the CCD chip causes significant fringe patterns. These fringe patterns can have very

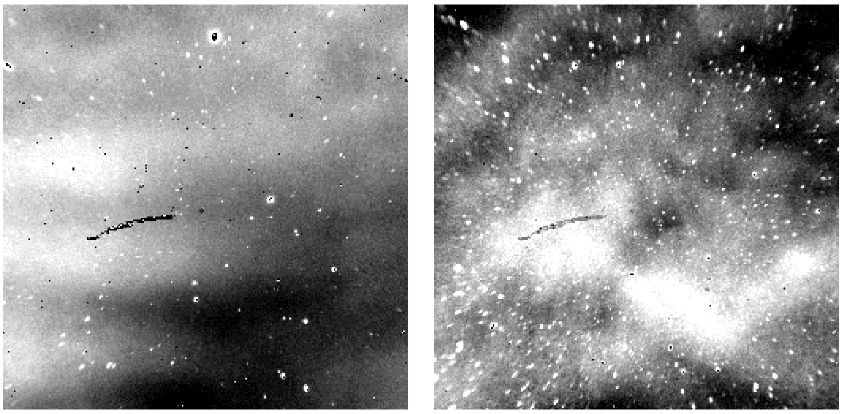

Figure 1: H- and J-band airglow images obtained with the 2MASS prototype camera by J. Adams & M. Skrutskie (1996). The field of view is about $10^\circ \times 10^\circ$.

complex geometries and topologies, and they occur on small scales down to a few pixels, comparable to object sizes.

Sky emission in the near-infrared wavebands (Y, J, H, K) is mainly due to vibrational-rotational transitions of OH, as well as O_2 and H_2O in the K band. The sky emission is bright, but more importantly it varies significantly over space and time, so it has to be carefully subtracted before NIR images can be used for photometry of astronomical objects.

As gravity waves pass through the upper atmosphere at an altitude of 80 to 105 km, both the strength and the structure of the sky emission vary over short time scales on the order of a few minutes to an hour. Spatially, variations are seen on angular scales of several degrees down to the arcmin level (Fig. 1), which means that significant variation can be seen in typical NIR imager fields of view. Impressive movies of the sky behaviour in the NIR bands obtained with the 2MASS prototype camera in 1996 can be viewed online.[1]

3 Requirements for sky subtraction modules

The smooth sky emission in *optical* images does not affect photometric measurements if the imaging data are restricted to single CCD chips. The slow spatial variations of the sky emission across objects can be captured by a bilinear fit to the local background in an annulus around the object so that subtraction of a sky model is not necessary. However, for mosaic images that are assembled from multiple exposures with single- or multiple-chip cameras the sky has to be subtracted from the individual images since otherwise steep gradients might appear in the background at chip boundaries. As long as object sizes in the images are small compared to the scale over which the sky varies, a sky model can be obtained from the image by applying some sort of low-pass filter or low-order polynomial fit.

As explained in the previous section, the sky emission in the near-infrared bands presents a number of characteristics that influence the observing strategy and that make it harder to correct for than in optical bands:

[1] `http://astsun.astro.virginia.edu/~mfs4n/2mass/airglow/airglow.html`

Table 1: Typical sky brightness (in mag/arcsec^2) at Paranal. UBVRI data are from Patat (2004), JHKL from Cuby et al. (2000)

Band	U	B	V	R	I	J	H	K_s	L
Magnitude	22.3	22.6	21.6	20.9	19.7	16.5	14.4	13.0	3.9

- The sky is very bright in the near-infrared (cf. Table 1), which makes it the dominant component of raw NIR images. In the reddest bands (K_s and L) thermal emission from the telescope becomes a further significant background component.
- The sky shows structure on angular scales that are smaller than the field sizes of typical NIR instruments and can become comparable to the sizes of large objects in the field.
- Sky emission is variable on time scales as short as five to ten minutes, affecting both the sky brightness and the spatial structure.

NIR observations typically consist of a large number of dithered exposures with integration times of a few seconds to avoid saturation by the sky background. This makes it possible to build a sky correction frame for each individual exposure by stacking a certain number of exposures taken within a time window such that the sky does not vary significantly. Due to the dithering, object flux can be effectively removed from the combined sky image. This basic strategy is implemented in all reduction pipelines for NIR data, although the details (in particular concerning object rejection) may vary.

Object rejection works only if the dithering shifts are larger than the largest object sizes in the field. If the observations target an object which covers a large fraction of the field, one typically observes a blank “sky field” some distance away from the target, alternating sky and object exposures in a certain pattern to achieve good temporal sampling of the sky behaviour. This *nodding strategy* is not different from the *dithering strategy* from an algorithmic point of view; an implementation of a sky subtraction recipe should however provide for the handling of both strategies by automatically identifying and using sky exposures if available.

The same strategies can be applied to the correction of *fringe patterns* in *I* or *z* band images.

A good sky subtraction delivers a corrected image that is as flat as possible “within the errors”. This means that the flux from all pixels (which do not contain flux from astronomical objects in addition to sky flux) should be consistent with being drawn from a noise distribution with a identical mean (the average sky flux level) and standard deviation as given e.g. by an error map. In addition to this test which treats pixels independently, one should look out for regions where the resulting background level is coherently high or low compared to the mean, although possibly at low significance for each individual pixel. Typically, this occurs around bright objects whose wings are not properly rejected in the construction of the sky correction. This leads to an oversubtraction of the sky which in turn might influence the photometry in this area.

The standard sky correction as sketched above produces a sky image which is noisy from pixel to pixel (Newberry 1999). Since this noise is introduced in the corrected images, care must be taken to keep its level as low as possible by using information from as many exposures as possible.

4 Other pipelines

Most major observatories operate a NIR imager, and there are a number of pipelines and software packages for the reduction of their data. In the following we mention several packages and describe the way they treat the sky background.

4.1 ‘I‘iwi

‘I‘iwi is the reduction pipeline for WIRCam, a $4 \times 2048 \times 2048$ mosaic camera on the Canada-France-Hawaii Telescope. The pipeline is described online[2] (incomplete), in a presentation at the CFHT User Meeting (2007)[3] and in Marmo (2007). It is written in IDL. To our knowledge it is not publically available, and we therefore could not test its performance.

Sky correction images are built from at least 3 adjacent observations. The images are normalized to 1 and objects are detected in each individual exposure using SExtractor (DETECT_THRESH = 1.5, ANALYSIS_THRESH = 1.5) and masked, where the mask is expanded (by applying a 5×5 smoothing kernel) to take into account object wings that have not been detected. The final correction image is the pixel-wise median over the exposures in the window.

4.2 ESO/MVM

The ESO/MVM pipeline was originally developed by Benoît Vandame and was used notably for the reduction of the ESO Imaging Survey (EIS). Development of ESO/MVM was discontinued at the end of 2006. The last version was 1.3.5.

The algorithms used in ESO/MVM are described in detail in Vandame (2004), sections 4.4 and 4.5. The correction for a smooth background from optical images uses a pyramidal median transform, an iterative smoothing with a median filter kernel of 3×3 pixels placed at every other input pixel. At every step, the image is thus reduced in size by a factor 2. The iteration stops at a last scale S which is the scale of the smoothly varying background. The filtered image at this step is referred to as the “mini-background”. It is resampled back to the original image grid by cubic spline interpolation and subtracted from the input image.

The pixel-based sky estimator used in MVM for NIR-background and fringe correction uses the usual two-pass technique, with object detection and masking from an intermediate (astrometrically) stacked image. A minor difference from both the methods in ‘I‘iwi and in our implementation is that MVM subtracts the mean from each exposure instead of dividing by it before the coaddition into the sky correction frame.

4.3 SExtractor, SWarp

`SExtractor` and `SWarp` form part of the Terapix software suite[4] along with several other programmes that are used for the reduction of wide-field images from MegaCam on the Canada-France-Hawaii Telescope but are of general use for the reduction and analysis of imaging data. `SExtractor` (Bertin & Arnouts 1996; Bertin 2009) is an object detection and analysis programme that is widely used in

[2] `http://www.cfht.hawaii.edu/Instruments/Imaging/WIRCam/IiwiVersion1Doc.html`
[3] `http://www.cfht.hawaii.edu/Instruments/Imaging/WIRCam/pics.WIRCam/iiwi_usersmeeting2007.pdf`
[4] Now called *astromatic*, `http://www.astromatic.net/`

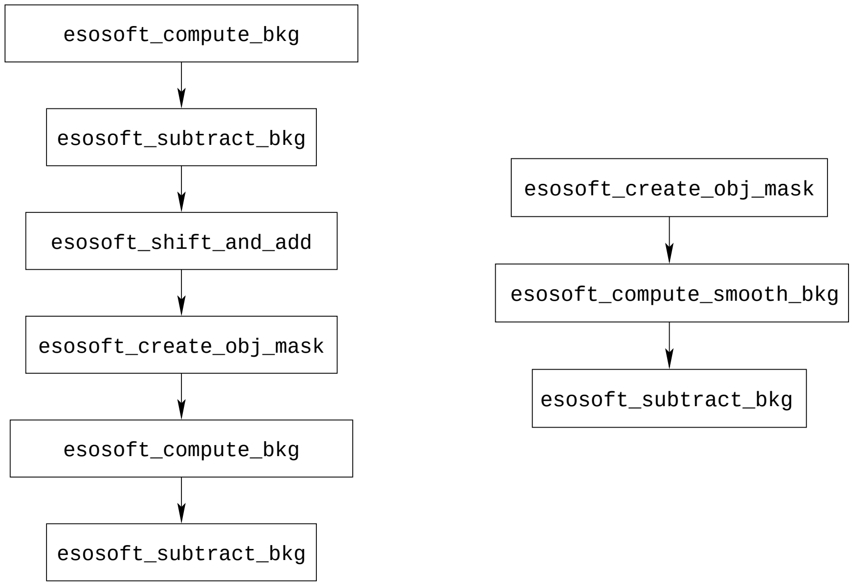

Figure 2: Sequence of recipes to be used for (left) the two-pass correction for NIR images and (right) the smooth correction for optical images.

the astronomical community. `SWarp` (Bertin 2008) resamples images onto a common output grid and homogenizes and stacks them. Both programmes contain a background estimation routine which is useful for estimating and removing smoothly varying backgrounds typical for optical imaging data.

The routine first establishes a grid over the image (with a default mesh size of 64 pixels) and computes an estimate of the background in each cell. For this purpose, a $\kappa\sigma$ clipping algorithm is used, along with estimation of the mode. After application of a median filter to suppress contamination from bright stars in some meshes, the grid is then interpolated with bicubic splines to result in the final background map.

The main parameter is the size of the grid meshes (`BACK_SIZE`). This should be larger than the largest objects in the field but smaller than the smallest scale over which the background changes appreciably.

5 Algorithm for NIR images: pixel-based sky estimation

We implement a two-pass strategy that constructs a sky correction image for each exposure I_i ($i = 1,\ldots,N_{\mathrm{exp}}$) from either a set of dithered object exposures $\{\mathsf{I}_j\}$ taken immediately before and after in the observing sequence or a sequence of sky exposures $\{\mathsf{J}_j\}$ taken in close temporal vicinity to the exposure that is to be corrected. Algorithmically, there is no difference between these two approaches and we will refer to the set of exposures used for the correction generically as $\{\mathsf{I}_i\}$. We will comment on the selection of this set in Sect. 9.2.

The basic assumption is that each exposure I_i can be represented as the sum of flux from astronomical objects T_i, a scaled version of a common sky structure S and noise ε:

$$\mathsf{I}_i = \mathsf{T}_i + c_i \mathsf{S} + \varepsilon\,. \tag{1}$$

The task of the sky subtraction is to estimate and isolate the scientifically interesting part T_i.

We estimate the scale c_i by the median μ_i of the image I_i. Each image can then be considered an estimator of S after division by its median if the object contribution T_i can be removed. To decrease the statistical noise, we compute the pixelwise average of images within a window:

$$\hat{\mathsf{S}}(x,y) = \frac{1}{N_{\text{good}}} \sum_i \frac{\mathsf{I}_i(x,y) - \mathsf{T}_i(x,y)}{\mu_i} \,. \tag{2}$$

In practice, the subtraction of T_i in Eq. (2) is done by identifying and discarding images from the sum in which the pixel (x,y) under consideration is affected by flux from an astronomical object. The number of remaining images for a given pixel (i.e. those that contain only sky flux at this position) is then $N_{\text{good}} \leq N_{\text{exp}}$.

The two-pass strategy is visualised in the left panel of Fig. 2. Its purpose is to reliably identify pixels that contain object flux by running an object detection algorithm on a preliminary (astrometrically correctly) combined image composed of exposures from the window I_j. The object catalogue is then used to mask objects from the individual exposures which are then averaged into the final sky correction image.

5.1 Stacking of sky-subtracted images

In order to create a more sophisticated mask based on application of an object-detection algorithm, the first-pass sky-corrected images are stacked into a combined image which allows detection of fainter objects than would be possible from individual exposures. The catalogue of detected objects is converted into a binary mask. The masked regions are generously expanded in order to robustly mask out even the faint wings of objects.

At this stage, distortion correction based on a precise astrometric solution is not necessary. We opt to align the exposures using simple integer pixel shifts, so that no rebinning of the images is required. The shifts are obtained from the headers of the input images.

In the case of Hawk-I, the telescope offsets applied during the execution of an observing block are stored in the fits header as `HIERARCH ESO SEQ CUMOFFSETX` and `HIERARCH ESO SEQ CUMOFFSETY`. Note that these are *telescope* offsets and therefore have opposite sign to the shifts of the image frames. The shifts Δx and Δy used here are already multiplied by -1 and are thus to be considered as *frame* offsets.

The situation is illustrated in Fig. 3. The size of the individual exposures is $n_x \times n_y$, the size of the stacked image encompassing all the exposures is $N_x \times N_y$. The position of the lower left corner of the reference frame within the stacked frame is then:

$$(1,1)_{\text{ref}} \longmapsto (1 - \min_i \Delta x_i \,,\, 1 - \min_i \Delta y_i)_{\text{stack}} \tag{3}$$

and the position of exposure i is

$$(1,1)_{\text{exp},i} \longmapsto (\Delta x_i + 1 - \min_i \Delta x_i \,,\, \Delta y_i + 1 - \min_i \Delta y_i)_{\text{stack}} \tag{4}$$

The size of the output image is

$$(N_x, N_y) = (n_x + \max_i \Delta x_i - \min_i \Delta x_i \,,\, n_y + \max_i \Delta y_i - \min_i \Delta y_i) \tag{5}$$

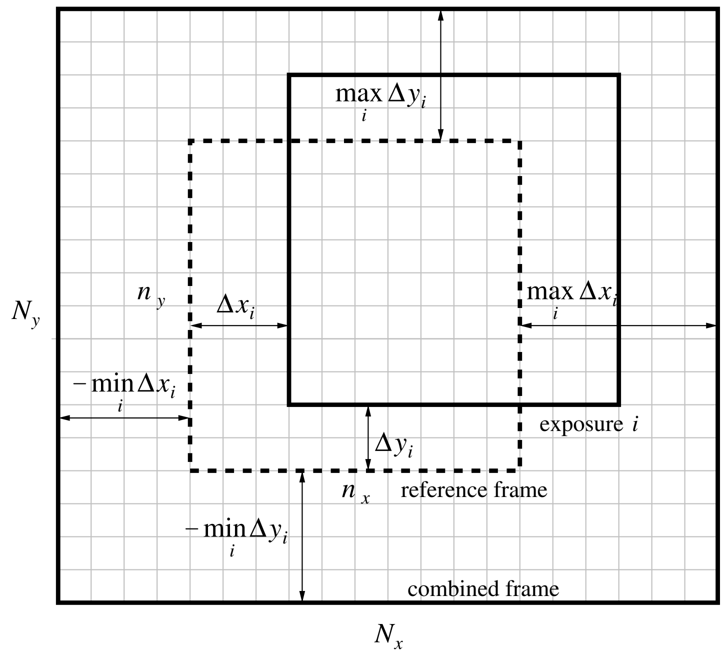

Figure 3: Position of an individual exposure within the stacked image frame. The shifts $(\Delta x_i, \Delta y_i)$ are given with respect to a reference position which is here represented by a corresponding reference frame.

To obtain the value of pixel (X,Y) in the stacked image, the average of the values of the input pixels shifted to that position is computed. For exposure i, the relevant pixel is

$$(x_i, y_i) = (X + \min_j \Delta x_j - \Delta x_i \, , \, Y + \min_j \Delta y_j - \Delta y_i) \tag{6}$$

5.2 Object detection and creation of object mask

The power of using an object detection algorithm to create a mask for rejecting object light lies in the fact that it detects *coherent* structures and is thus able to identify wings of objects where the flux excess due to the object is so low as to not be significant with respect to the background noise if measured on a pixel basis, yet high enough on average to lead to an overestimate of the estimated sky background.

Object detection algorithms (such as implemented in SExtractor) start by thresholding an image, where the threshold is typically given as a certain number of standard deviations of the background noise. The pixels above the threshold are then associated to objects using a filter, a friends-of-friends algorithm or fitting of a Gaussian profile.

For the purpose of creating an object mask, we first apply a morphological opening using a 3×3 square structuring element and then a growing kernel to the mask. Morphological opening removes structures from the mask that are smaller than the structuring element; it is used here to remove single-pixel detections which are mostly extreme values of the normal background noise distribution.

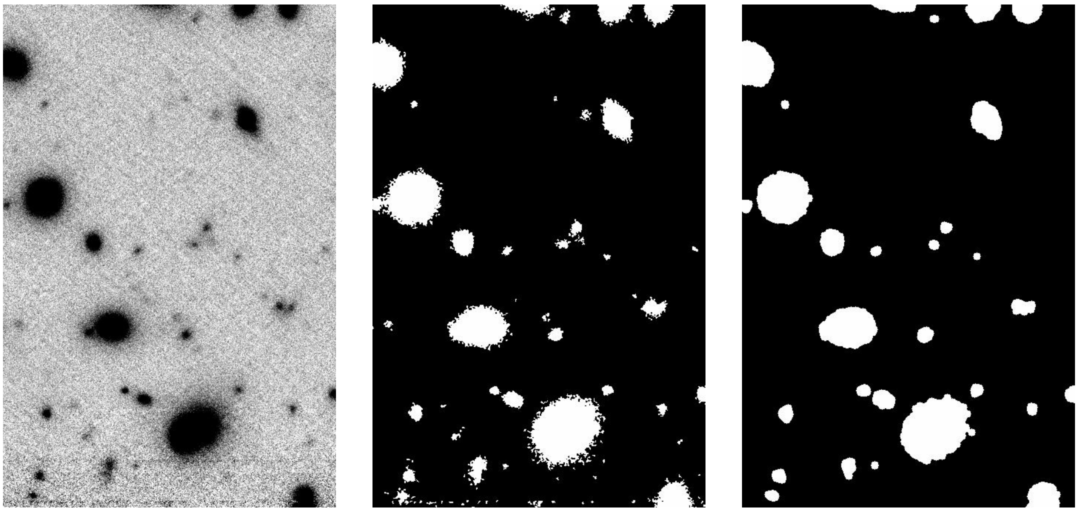

Figure 4: Demonstration of mask creation. Left: stacked image of 15 exposures of Abell 1689, HAWK-I, J band. Centre: SExtractor `SEGMENTATION` image. Right: EsoSoft mask after growing by 5 pixels. Both the SExtractor and the EsoSoft mask were created with threshold 1.5σ.

The growing kernel is a circular function of the form

$$k(x,y) = 1 - \frac{\sqrt{x^2+y^2}}{r_g} \tag{7}$$

within the growing radius r_g. Application of the kernel expands the masked region, extending further into the low surface-brightness wings of objects.

Fig. 4 illustrates the application of the mask creation algorithm and compares the result to a SEGMENTATION check image created by SExtractor.

5.3 Treatment of masks

Only pixel values that are representative of the sky flux at their position should be used in the computation of the sky background image. Objects will be masked out using *object masks*, created with an object-detection algorithm as described in Sect. 5.2. The recipe should also handle static and individual bad-pixel masks and, if necessary, update them. Pixels in the sky image that will be flagged are those where the number of "useful" input values application of bad-pixel and object masks and outlier rejection is less than a specified minimum number.

Hence, the recipe will take on input N images; 0, 1 or N bad-pixel masks (a single mask will be applied to all images); 0 or N object masks. On output there will be 0 or N updated masks.

5.4 Noise propagation

It is assumed that the input images are accompanied by noise images which we take to contain the *variance* σ^2 at each pixel position. It is also assumed that the noise is independent from one pixel to the next. If necessary, the noise distribution is taken to be Gaussian, which is appropriate if the data are photon noise dominated. Since photon detection is really a Poisson process, the variance of the asymptotic Gaussian distribution of photon (or electron) counts is $\sigma^2 = \mu$. For an image given in analog data units (ADU), the gain factor g in ADU/e$^-$ has to be taken into account, giving $\sigma^2_{\mathrm{ADU}} = g^2\mu$.

The background correction images created in the two-pass method are themselves noisy and their subtraction from the input images will lead to increased noise in the corrected output images. We attempt to take the propagation of noise properties through the recipes into account by outputting updated weight maps along with the data products.

The computation of a pixel value in the background image is essentially an estimate of the average of a sample of N pixel values from the exposures in the set. Each of these is taken from an independent noise distribution $P(\mu_i, \sigma_i, \dots)$, with $i = 1, \dots, N$. The variances σ_i^2 can vary from exposure to exposure, e.g. due to variable sky background or varying exposure time. We are therefore in general dealing with *heteroscedastic* data.

The noise variance in the sky correction image depends on the method used to combine the pixel values, i.e. the estimator of the mean of the sample. The propagation of the noise is discussed in the presentation of the different methods in Sect. 7.

The variance in the corrected image is simply the sum of the variances in the input image and the sky correction image:

$$\sigma^2(\mathrm{BKG\_CORRECTED}) = \sigma^2(\mathrm{BASIC\_CALIBRATED}) + \sigma^2(\mathrm{BKG\_IMG}). \tag{8}$$

The noise in the correction image does not correlate with the noise in the image to be corrected since the latter is explicitely not included in the construction of the correction image. As a rough estimate, the variance in the correction image will be $\sim 1/N_{\mathrm{images}}$ times the variance in the input image.

6 Algorithm for optical images: smooth sky estimation

Images taken in the visible wavebands usually show a background which varies slowly over the size of the detector. In cases where objects are fairly sparse and small so that the majority of pixels in the image carries pure sky signal it is possible and customary to obtain a smooth estimate of the sky background from the exposure to be corrected. This estimate can be obtained in a parametric way, say by fitting a polynomial or spline function, or non-parametrically by heavily smoothing the image with a large kernel or reconstructing a background estimate on a grid with a mesh size that is larger than the largest objects in the field. The latter approach is implemented in `SExtractor` and `SWarp` (see Sect. 4.3).

Here, a fit using orthogonal Legendre polynomials in pixel coordinates x and y is used. The method is implemented in the recipe `esosoft_compute_smooth_bkg`, which uses functions from `esosoft_linalg.c`, `esosoft_matrix.c` and `esosoft_orth_poly.c`. The latter extend the functionality of CPL for general linear algebra, matrix computations and orthogonal polynomials.

The procedure for the background correction for optical images is as follows (cf. the right hand panel in Fig. 2):

1. detect objects and create an object mask (`esosoft_create_obj_mask`)
2. fit a smooth background model (`esosoft_compute_smooth_bkg`)
3. subtract the background model (`esosoft_subtract_bkg`)

The preparatory step attempts to clean the set of pixels over which the polynomial is fitted from object contamination. The method and the recipe are the same as described in Sect. 5.3. In contrast to the

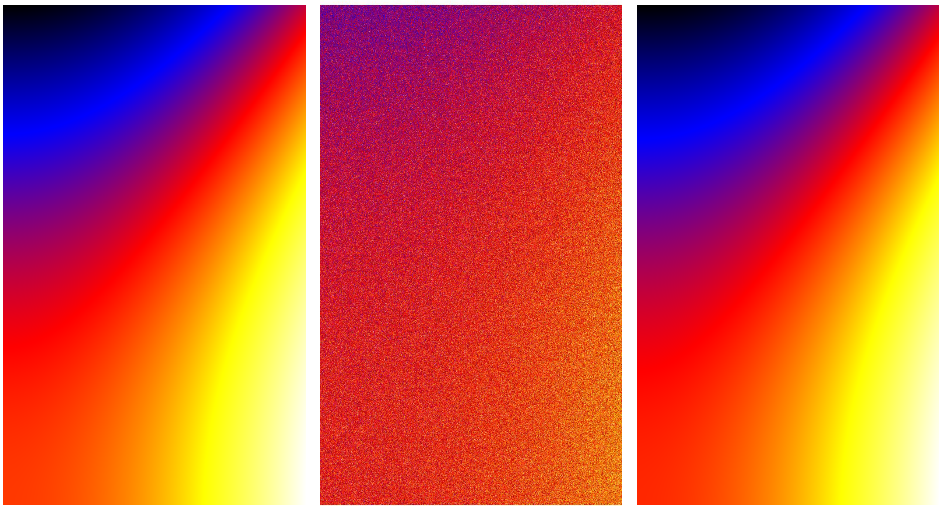

(a) The original image. (b) The original image with added noise. (c) The regularised noisy image.

Figure 5: Least squares regularisation of the sky background.

NIR images, the mask is produced on the individual exposure instead of a stacked image. This is possible because the fraction of pixels contaminated by possibly undetected faint objects is expected to be low compared to the number of pixels with pure sky so that the smooth fit will be unaffected by their presence.

Next, the noisy sky background is regularised, and partially denoised, by means of a least squares fit with the polynomials of two variables.

The least squares fit is computed using the Wiener filter, the method is also known as the Tikhonov regularisation or the ridge regression (e.g. Björck 1996). The method requires the solution of the normal equations with the diagonal incremented by a constant reflecting the conditioning of the problem and the amount of noise (e.g. the variance of the Gaussian noise).

The resulting polynomial approximation is a linear combination of the basis polynomials. For reasons of numerical stability, the basis consists of the tensor products of the Legendre polynomials. Specifically, each basis polynomial has the form $P_i(x)P_j(y)$, where P_i is the Legendre polynomial of degree i rescaled according to the size of the image. When used with rectangular images, the Legendre polynomials in the variable x have a different scaling than those in the variable y. The most practical basis consists of all tensor products $P_i(x)P_j(y)$ with $i+j \leqslant k$, where k is a fixed integer.

As an example, Fig. 5 presents regularisation results obtained with 6 tensor products at the 15% noise level.

7 Robust estimators of location – outlier rejection

Stacking of images and spectra by averaging or summing is a standard procedure in astronomical data reduction and is part of several deliverables in the context of the ESO Data Reduction Software Project. Astronomical data are often affected by cosmic ray hits, satellite or airplane trails, reflections etc., i.e. data values which do not arise from the usual noise distribution of the astronomical image and which should not appear in the stacked image. In the course of the two-pass method for NIR sky subtraction, images are stacked in the image frame with the goal of removing light from stars and other astronomical sources. In the list of pixel values to be combined all of these effects show up as outliers from the underlying noise distribution. Rejection of outliers is thus a requirement for a satisfactory stacking module. We discuss here a number of robust estimators of the mean of a sample of data values and outlier rejection methods which have been implemented in the EsoSoft modules.

7.1 Arithmetic mean

The arithmetic mean of a sample $x_1, x_2, \ldots, x_N$ is defined as usual as

$$\bar{x} = \frac{1}{N} \sum_{i=1}^{N} x_i \,. \tag{9}$$

For a random variable X that is normally distributed ($N(\mu, \sigma)$), the sample mean is the most efficient estimator of the population mean μ. By the central limit theorem this statement holds asymptotically for many noise distributions encountered in practice.

If the noise distribution of X is unknown, its variance can be estimated from the sample as

$$\hat{\sigma}^2 = \frac{1}{N-1} \sum_{i=1}^{N} (x_i - \bar{x})^2 \,. \tag{10}$$

If X has variance σ then the average (9) has variance

$$\sigma_{\mathrm{av}}^2 = \frac{\sigma^2}{N}. \tag{11}$$

This relation can be used in place of Eq. (10) if an *a priori* estimate of the noise in the images is available and is the same for all images (homescedastic case).

When each image has its own noise image $\sigma_i^2(x,y)$ (heteroscedastic case), then the weighted mean

$$\bar{x} = \left(\sum \frac{1}{\sigma_i^2} \right)^{-1} \left(\sum \frac{1}{\sigma_i^2} x_i \right) \tag{12}$$

can be used, giving higher weight to images with lower noise. The variance of the average value at a given pixel is in this case given by

$$\frac{1}{\sigma_{\mathrm{av}}^2} = \sum_{i=1}^{N} \frac{1}{\sigma_i^2} \,. \tag{13}$$

While the average is efficient, it is also sensitive to the presence of outliers, i.e. sample values which arise from a process which is not included in the $N(\mu, \sigma)$ model of the background noise. Cosmic ray hits and the like that affect single exposures in the set therefore show up clearly in the stacked image. The average fails in the presence of a single outlier in the sample.

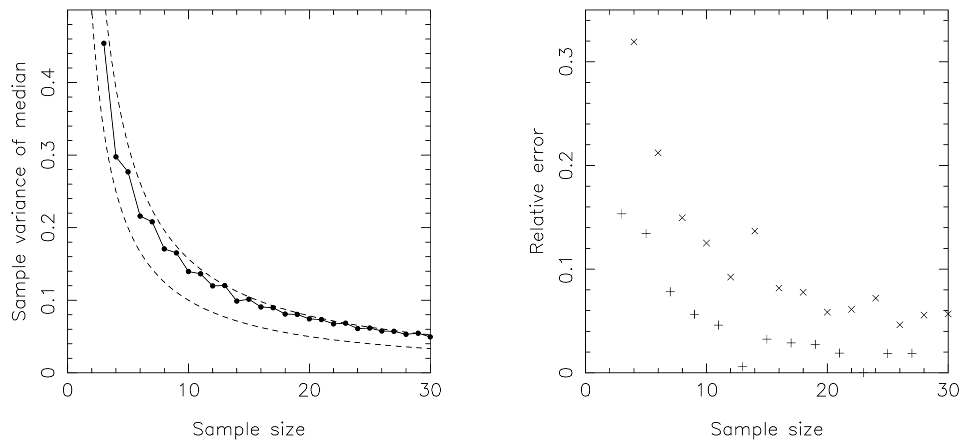

Figure 6: Sample variance of the median as a function of sample size (left panel). The lower dashed curve is the expected variance of the arithmetic mean, the upper dashed curve the asymptotic variance of the median. The right panel shows the relatve error compared to the asymptotic variance. Crosses are for even sample sizes, pluses for odd sample sizes.

7.2 Median

For the computation of the sample median, the sample $x_1, x_2, \ldots, x_N$ is first sorted such that $x_1 < x_2 < \cdots < x_N$. The median is then the 50th percentile of the ordered set, i.e.

$$x_{\text{med}} = x_{\frac{N+1}{2}} \quad \text{if } N \text{ is odd} \tag{14}$$

and

$$x_{\text{med}} = \frac{1}{2}(x_{\frac{N}{2}} + x_{\frac{N}{2}+1}) \quad \text{if } N \text{ is even.} \tag{15}$$

This definition for the median of even-sized samples ensures that the median is an unbiased estimator of the population mean μ for a Gaussian (in general symmetric) error distribution.

The sample median has a larger variance than the sample average, i.e. it is less efficient as an estimator of the population mean. It is however robust against outliers in the sample since only the cenral one or two sample values are actually used for the computation. The median only fails if half or more of the sample values are outliers, i.e. if a given pixel is affected by a cosmic ray hit in half or more of the images in the stacking list.

While the sample distribution of the median (which is not Gaussian) can be written down easily, the computation of its moments and hence its variance is difficult. A scheme to compute it analytically has been described by Maritz & Jarrett (1978) and exact values for a few sample sizes and a number of parent distributions are given in Rider (1960). However, these are hardly useful for practical purposes. The asymptotic value for the ratio of the variances of mean and median (the *asymptotic relative efficiency* of the median, DasGupta 2008) is

$$\frac{\sigma^2(\text{Mean})}{\sigma^2(\text{Median})} = \frac{2}{\pi}. \tag{16}$$

For a Gaussian distribution, $\sigma^2(\text{Mean}) = 1/N$. In Fig. 6, we determine the variance of the median from simulations by drawing for any given sample size 10 000 samples from a standard normal distribution.

The asymptotic value is approached from below, which makes it a conservative choice for an estimate of the median's variance. The relative error, defined as

$$\frac{\pi/2N - \hat{\sigma}^2(\text{Median})}{\pi/2N}, \tag{17}$$

is plotted in the right hand panel of Fig. 6. For even sample sizes, the relative error is significantly larger than for odd sizes because the variance is actually lower. For odd-sized samples, the error is below 10% for $N \geq 7$ ($N \geq 5$ according to Rider 1960), whereas for even sized samples this threshold is not crossed until $N \gtrsim 12$. However, as mentioned above, the asymptotic value in Eq. (16) is conservative for any N and we choose to use it throughout.

7.3 Min-max rejection

In the min-max rejection algorithm, the N_{low} lowest and N_{high} highest values are removed from the set and the average of the remainder is computed. For the ordered set x_i, the estimator is thus

$$x_{\text{min-max}} = \frac{1}{N - N_{\text{low}} - N_{\text{high}}} \sum_{i=N_{\text{low}}+1}^{N-N_{\text{high}}} x_i \,. \tag{18}$$

The median can be seen as the extreme case of min-max rejection where all but the central one or two data values are rejected. The variance of the estimator defined in Eq. (18) is hard to compute exactly and we have not been able to find an analytic formula for it. Intuitively, one expects the variance of the min-max rejection estimator to lie between those for the arithmetic mean of the full sample and the median; this is confirmed by the simulation shown in Fig. 7. 10,000 samples of size $N = 21$ were drawn from a standard normal distribution and submitted to the min-max rejection algorithm for $n_{\text{low}} = n_{\text{high}} = 0, 1, \ldots, (N-1)/2$. The first value is equal to the arithmetic mean, the last value is equal to the median.

7.4 Kappa-sigma clipping

In the min-max rejection algorithm, a fixed number of sample values is rejected, regardless whether they can be identified as outliers or not. If one wants to retain all "good" sample values and only reject true outliers, one has to employ an adaptive method which compares each sample value to the distribution of the entire sample.

In the $\kappa\sigma$ clipping algorithm, all values that deviate from the mean by more than κ standard deviations are rejected as outliers. Typically, $\kappa = 3$. The mean is usually estimated from the data, as is the standard deviation if no independent error estimate is available. One can introduce an iteration which stops once no more values are rejected.

We have found that standard $\kappa\sigma$ clipping is not very efficient in removing low-significance outliers and explain this by the fact that the estimates of the mean and the standard deviation are actually affected by the outliers themselves. Using more robust estimators of location and scale, like the median and the inter-quartile range, improves the situation and makes $\kappa\sigma$ clipping a robust and easy-to-use adaptive outlier-rejection algorithm.

$\kappa\sigma$ clipping has one parameter, the clipping threshold κ. The value of κ is not critical if outliers are expected to lie far from the sample distribution, as is the case for cosmic ray hits. Weaker effects, such as airplane trails, may require fine-tuning of κ if the method is able to capture these outliers at all.

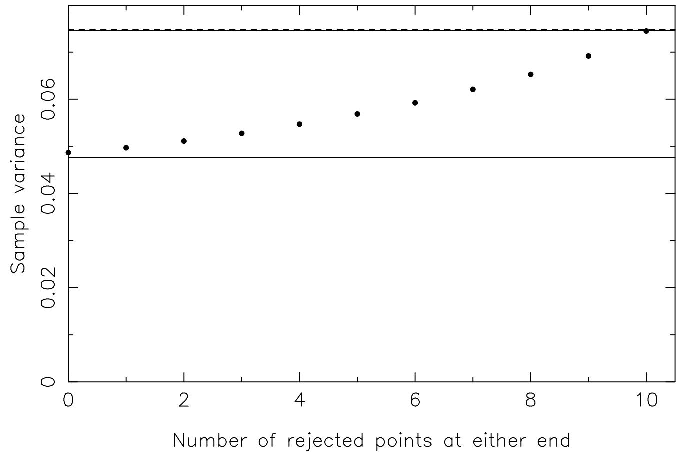

Figure 7: Behaviour of the sample variance of the min-max rejection estimator for a sample with $N = 21$ drawn from a normal distribution with $\mu = 0$, $\sigma = 1$ as a function of the number rejected data values $n_{\text{low}} = n_{\text{high}}$. The lower horizontal line give the expectation for the average, $\hat{\sigma}^2 = 1/N$, the upper solid line the measured variance for the median (equal to the point for 10 rejected points at either end), and the upper dashed curve the asymptotic expected variance for the median, $\frac{\pi}{2}\frac{1}{N}$.

The method might fail if more than a quarter of the input images are affected by outliers at a given pixel because then the estimate of the inter-quartile range might be affected and cause outlier rejection to fail. The method also requires a sufficient number of input images to be able to obtain reasonably accurate estimates of the mean and standard deviation.

The variance of the estimator now varies from pixel to pixel depending on the number of values that are rejected as outliers. Since outliers are not drawn from the same distribution as the “good” data values, the cleaned sample is equivalent to a smaller sample drawn from the noise distribution. In the general heteroscedastic case, the variance is therefore estimated as

$$\sigma_{\kappa\sigma}^2 = \frac{1}{N_{\text{good}}^2} \sum_{\text{good}} \sigma_i^2 . \tag{19}$$

7.5 Mean-median ratio

We introduce another adaptive outlier rejection method that exploits the different robustness of the mean and the median. It will be seen that the method suffers from difficulties in setting the parameters that make a general use of the method problematic. Its inclusion here should be seen as experimental.

Both the (arithmetic) mean and the median are estimators of the expectation value of a random variable X based on a sample x_i. In the presence of outliers, the mean reacts very sensitively whereas the median is hardly affected as long as the number of outliers is less than half the sample size. On the other hand, the mean is the most efficient and hence the preferred estimator of the expectation value if the distribution of X is Gaussian.

We make use of the different robustness of the mean and the median in order to test for the presence of outliers in a given sample. The test statistic is their ratio, whose sample distribution can in principle

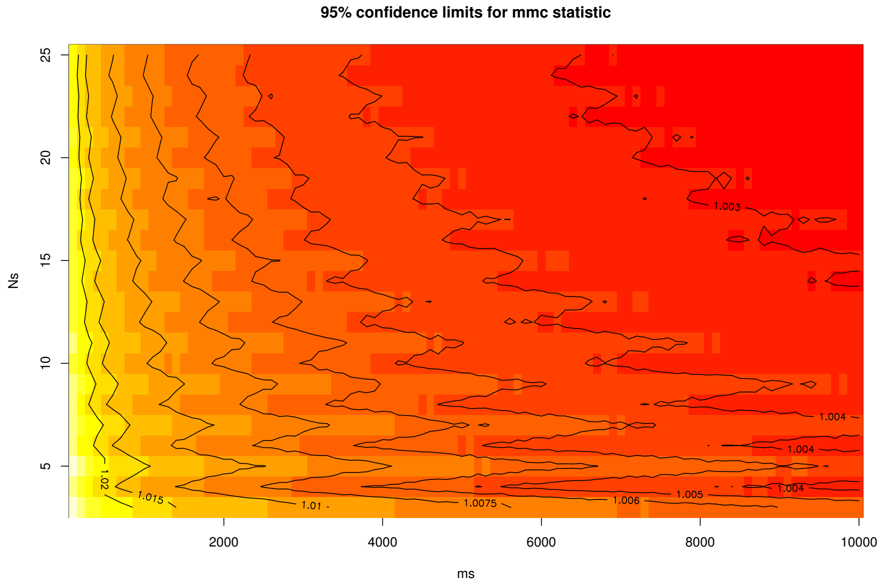

Figure 8: 95% confidence limits (single-tailed) for the mean-median ratio, as a function of population mean (variance is assumed to be equal to the mean) and sample size. The limits are derived from simulations, with 10000 samples drawn for each combination of m_s and N_s.

be computed under the null hypothesis H_0: "The sample is drawn from a Gaussian distribution with mean μ and dispersion σ." The expectation of the ratio is 1, hence if the value of the ratio deviates significantly from 1, the null hypothesis is rejected and it is assumed that outliers are present. In this case, we remove the highest (lowest) value from the sample if the ratio is larger (smaller) than 1. The ratio is recomputed on the cleaned sample and the procedure repeated until no more values are rejected.

Unfortunately, the sample distribution and the, say, 95% confidence limits for the mean-median ratio depend on the sample size (and due to the implication of the median it matters whether the sample size is odd or even) and the population mean and dispersion. While in principle it should be possible to compute the sample distribution analytically, we have not been able to do so yet. Fig. 8 shows results from simulations demonstrating the complexity of the confidence limit.

Testing on (Hawk-I) data, we have found that the performance over most of the field does not actually depend sensitively on the adopted threshold. The advantage of the mean-median ratio over min-max rejection is that it is adaptive: data values are only rejected if there is evidence that outliers are present. From our tests (see Sect. 9.1), the mean-median ratio appears to be more stable than $\kappa\sigma$ clipping when the latter is using the usual non-robust estimators of location and scale. However, the $\kappa\sigma$ method works better if the median and inter-quartile range are used as discussed in Sect. 7.4. The usefulness of the mean-median ratio until a way to automatically adapt the threshold to the data is found. Indeed, the method is likely to fail around bright objects if only a single threshold is used for the entire image, adapted to the general background level as demonstrated in Sect. 9.1.

8 Description of the code

8.1 Requirements

The recipes are pluggable data reduction modules (PDRM) as described in Sect. 3.5 of the CPL User Manual. Both the recipes and the functions are based on the Common Pipeline Library. We have linked and tested our code against CPL versions 5.0.1 and 5.2.0, We do not expect the code to work with older versions of CPL and do not support them.

The CPL library has to be linked against the `cfitsio` library[5]. We use versions 3.09 and 3.24 of this library. The former is used in ESO's VLT software and is recommended in the CPL User Manual. The current version, 3.25, was released on 9 June 2010 but has not been tested by us yet. We have not found any problems linking to version 3.24.

We also require the CPL library to be linked against the `wcslib` library[6] which handles astronomical world coordinate systems (WCS) in fits files. The EsoSoft recipes were tested with version 4.4.4 of `wcslib`. The latest version, 4.5.1, does not contain changes to the API and is expected to work as well.

8.2 Installation

The recipes and functions are delivered as a gzip'd tar ball. The current prototype release containing code for DR01 is `esosoft-0.0.3.tar.gz`, delivered on 3 September 2010. Unpacking the tar ball with

```
tar xvzf esosoft-0.0.3.tar.gz
```

creates a new directory, `esosoft-0.0.3`. To compile and install the recipes, run

```
configure --prefix=/path/to/pipelines [--with-cpl=/path/to/cpl/installation]
make
make install
```

The path to the cpl installation directory is only necessary if the library is not installed at a standard location. After installation, the recipes will be found in the directory

```
/path/to/pipelines/lib/esosoft/plugins/esosoft-0.0.3/
```

If `--prefix` is not specified, the default directory `/usr/local` is used. A separate directory for CPL recipes will, however, be found to be more convenient by most users.

To work with the recipes it is necessary to `esorex` to this directory with the command line option `--recipe-dir`, the parameter `esorex.caller.recipe-dir` in the configuration file `.esorex/esorex.rc` or the environment variable `ESOREX_PLUGIN_DIR`.

The code is documented with `doxygen` directives. To obtain the reference documentation, run

```
doxygen Doxyfile
```

To browse the html version of the manual, point your browser (e.g. `firefox` or `lynx`) to `html/index.html`. The directory `html` can be moved to any convenient location.

A hyperlinked pdf version of the manual can be obtained with the commands

```
cd latex/
make
```

This results in `refman.pdf`.

[5] `http://heasarc.gsfc.nasa.gov/fitsio/`

[6] `http://www.atnf.csiro.au/people/mcalabre/WCS/`

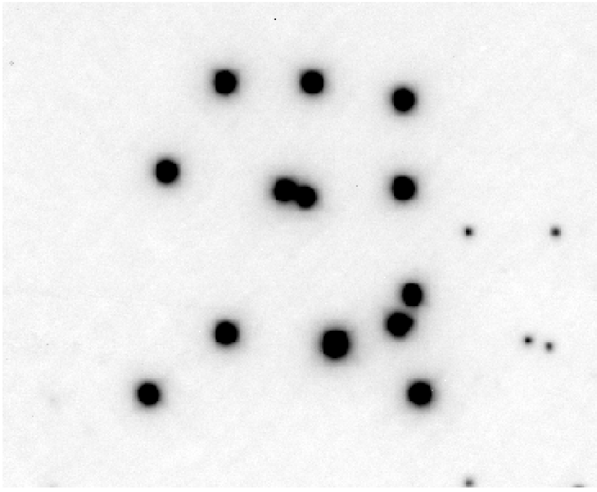

Figure 9: Dither pattern of the 15 Hawk-I exposures of Abell 1689. A section of chip 1 centred on a bright star is shown.

9 Application to test data

9.1 HAWK-I: Abell 1689, J band

This dataset contains observations of the cluster of galaxies Abell 1689, obtained with Hawk-I in the *J* band. There are 15 exposures taken within a single OB on 2008-05-03.

The field is on average not very crowded (2.4 % of pixels are covered by objects[7]), the cluster is located in the center of the field at the cross between the four chips. At redshift $z = 0.185$, the largest galaxies have diameters of about 7 arcsec, i.e 60 pixels. The maximum dithering offsets used in the observations were about 275 pixels, significantly larger than the size of the largest objects, hence the object images themselves can be used to create a sky correction image.

There are, however, two pairs of offset positions which are very close together, such that the images of bright stars from these exposures overlap (Fig. 9).

We compare six images, background corrected with different parameters. Except for image 1, all images were treated with the two-pass method with the same object mask. The exposure chosen was number 7 in the set; with the half window set to 7 exposures (the default), the background correction for this exposure was created from $6+7 = 13$ exposures.

- Image 1: median-stacked, first pass only, i.e. no object masking was done
- Image 2: average-stacked (created using min-max rejection of zero outliers)
- Image 3: min-max rejection of the three highest and the three lowest data points
- Image 4: outlier rejection by the mean-median ratio, with threshold set to 1.001
- Image 5: outlier rejection by kappa-sigma clipping, with $\kappa_{\text{low}} = \kappa_{\text{high}} = 3$
- Image 6: median-stack, after second pass

[7]based on a SExtractor run on a sky-corrected individual exposure with `DETECT_THRESH` = 1.5. The mean value of the binarized `SEGMENTATION` check image gives the object cover fraction.

Table 2: Statistics on empty rectangular patch in background-corrected exposure 7 of the A1689 set. IQR is the inter-quartile range.

No.	Description	Mean	Median	Std. dev.	Lower quartile	Upper quartile	IQR/1.349
1	med_1st	−1.31647	−1.28027	38.0784	−26.6416	23.7188	37.33165
2	mm_rej_0	−1.21542	−0.64404	36.7886	−26.0146	24.4199	37.38658
3	mm_rej_3	−1.12121	−0.54053	37.3904	−25.7520	24.2871	37.09348
4	mm_ratio	+0.17645	+0.36914	37.2070	−25.0059	25.3877	37.35626
5	ksigma	−1.12455	−0.96680	37.7574	−26.5225	24.0078	37.45760
6	med_2nd	−0.97987	−0.83594	38.0031	−26.1553	24.1055	37.25782

Table 2 shows the statistics on a rectangular patch of 239×239 pixels in chip 1 which is free from detectable objects.

The stacking method has only a small influence on the statistics of the product. The trends are as expected: The straight average (image 2) has the smallest noise, the median (images 1 and 6) the highest. The adaptive and non-adaptive outlier rejection methods have intermediate noise. However, the differences are small.

The performance of the outlier rejection is illustrated in Fig. 10 where a section of the image near a bright star is shown. The large scale correction is fine for all methods, differences can be seen on the scale of individual pixels. In the image section, the median stacks, the min-max rejection with `rejlow = rejhigh = 3` and the kappa-sigma clipping (images 1, 3, 5 and 6) produce corrected images that are free of artifacts, indicating that the outlier rejection worked satisfactorily. The straight average (image 2) shows artifacts near the bright star; this is expected because the average is not a robust estimator of location, i.e. it is sensitive to outliers. Also, the mean-median ratio method (image 4) shows artifacts near the bright star. While the method performs well over most of the image, it is unfortunately the case that the threshold for rejection depends on a number of parameters of the image, for instance the noise level in the image. In the wings of the star, the noise is higher than the average background noise, so that the threshold should be adjusted, otherwise the method might easily fail.

The choice of default method for the creation of the background correction is based on these results. The choice is a trade-off between *statistical efficiency* and *robustness*. For the default method, it is argued that robustness should be the prime criterion: The default method should produce clean and useable results for the majority of the expected data sets without requiring the user to adjust default parameters. We therefore choose the *median* as the default method, the most robust method amongst those available for `esosoft_create_bkg`. The median does not have any adjustable parameters and can be applied "blindly" to any data set. However, as shown above, the median creates corrected images which are somewhat noisier than the images created with the other methods. Users who want to improve the noise characteristics of their products will have to use one of the other methods. For simplicity, the min-max rejection can be used with the number of rejected pixels adapted to the number of exposures in the set or sliding window[8]. Kappa-sigma clipping (using robust estimators for location and scale) is an adaptive rejection method which works reliably in many situations. We provide the mean-median ratio constraint for completeness but point out the difficult question of how to set the threshold.

[8]In fact, the median can be viewed as the extreme case of min-max rejection where all but one or two pixel values are rejected.

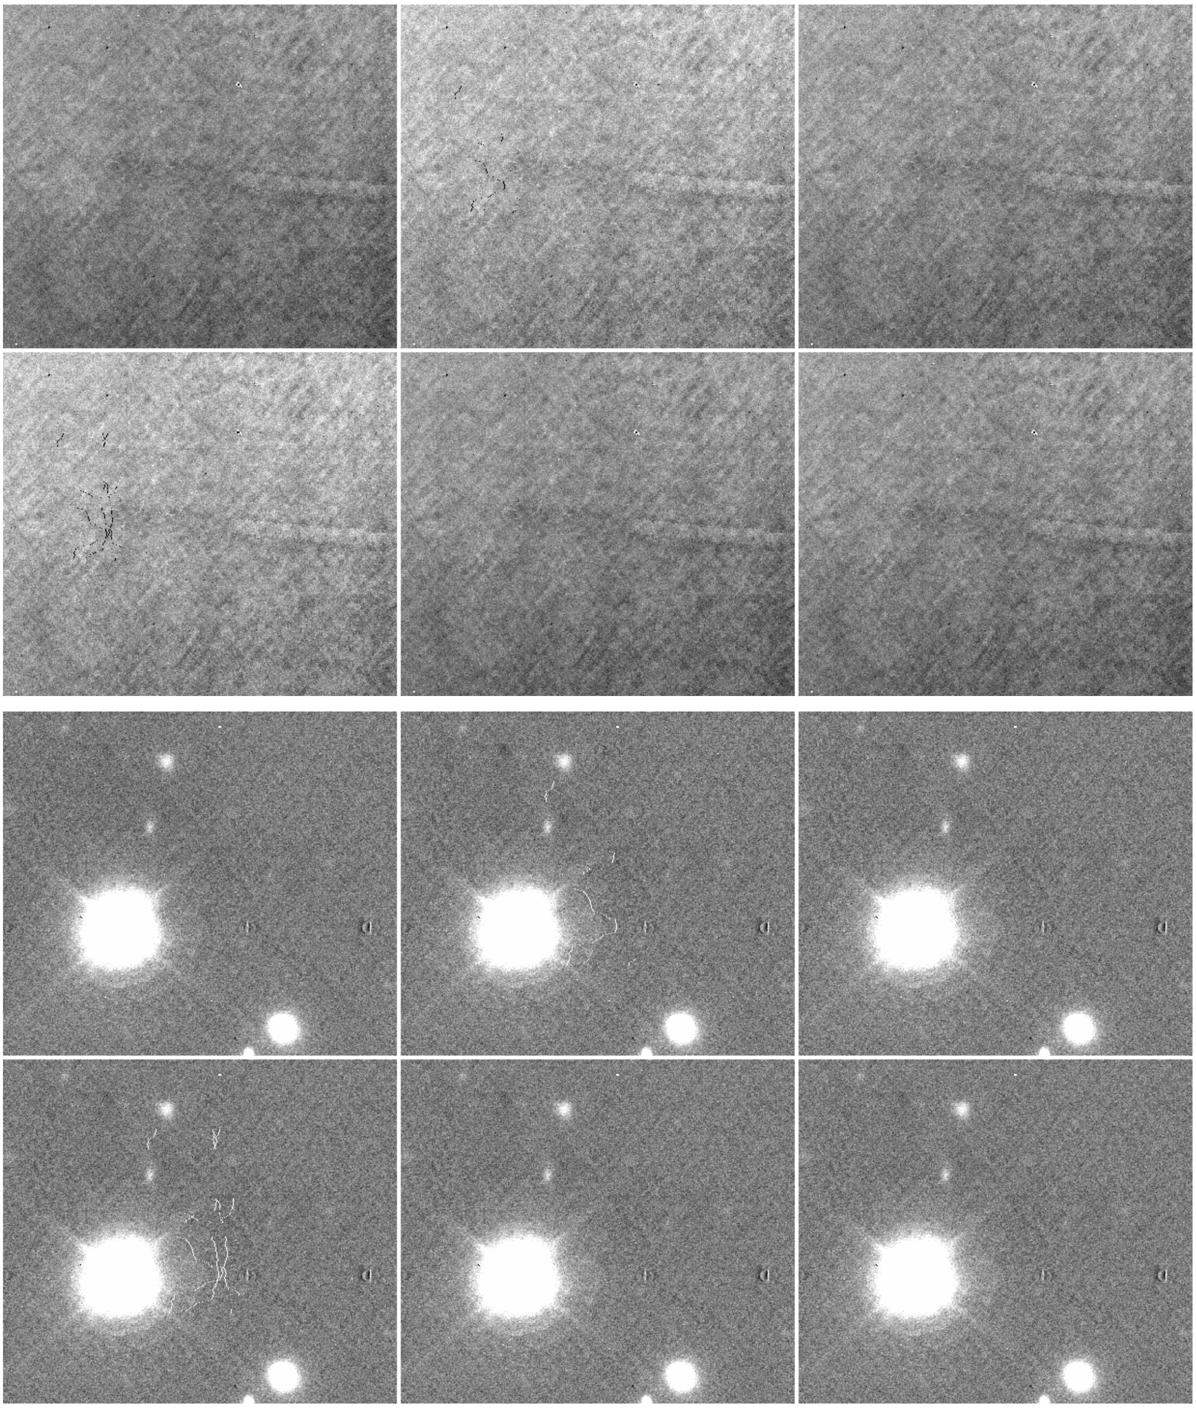

Figure 10: The top six panels show background correction images for an exposure from the A1689 HAWK-I set created using the following methods (images are numbered rowwise left to right): (a) median stacking, no object masking; (b) arithmetic mean; (c) min-max rejection with `rejlow = rejhigh = 3`; (d) mean-median ratio; (e) $\kappa\sigma$-clipping; (f) median with object masking. The bottom panels show a section of the exposure after subtraction of these correction images. Artifacts related to the sky subtraction are visible in Image (b) and (d). The parameters for these images are the same as in the corresponding image in the upper two rows.

9.2 HAWK-I: NGC 7793, K band

NGC 7793 is a large spiral galaxy with a diameter of about 7 arcmin which fills the greater part of the field of view of Hawk-I ($7.5 \times 7.5\,\mathrm{arcmin}^2$) and which makes it impossible to determine the sky contribution from the object images themselves. The observers of the data set used here took this problem into account by taking exposures of a blank patch of sky, located about 15 arcmin north of NGC 7793, inbetween the object exposures.

The exposures used here come from an OB executed 2008-11-10 using Hawk-I in the *K* band. There are 26 object exposures and 6 sky exposures taken in the following sequence:

OOOOOOOOOOO SS OO SS OO SS OOOOOOOOOOO

As there are only few sky exposures concentrated in the middle of the OB and none of the object exposures are useable for determining the sky correction, the best solution is to use all of the sky exposures to construct a single sky correction image which is applied to all the object exposures from the OB.

The EsoSoft background subtraction recipes do not explicitely provide for using separate sky exposures. This can, however, easily be achieved by writing appropriate SOF files.

In the present case, the SOF file for `esosoft_compute_bkg` contains all the "sky" files. In addition, the first sky file is entered twice:

hawki_step_basic_calib_sky001.fits BASIC_CALIBRATED
hawki_step_basic_calib_sky001.fits BASIC_CALIBRATED
hawki_step_basic_calib_sky002.fits BASIC_CALIBRATED
hawki_step_basic_calib_sky003.fits BASIC_CALIBRATED
...

The first output image, `esosoft_sky_bkg_001.fits`, will be combined from all six sky exposures and will be the one used to correct the object images. The SOF file for `esosoft_subtract_bkg` then reads:

hawki_step_basic_calib_obj001.fits BASIC_CALIBRATED
esosoft_sky_bkg_001.fits BKG_IM
hawki_step_basic_calib_obj002.fits BASIC_CALIBRATED
esosoft_sky_bkg_001.fits BKG_IM
...

Whether a second pass is necessary with this approach depends on the density of objects and the size of the dither shifts between the sky exposures.

9.3 HAWK-I: IC 1613, K band

IC 1613 is an irregular dwarf galaxy which covers and actually exceeds the field of view of Hawk-I. Unlike NGC 7793, it is resolved into individual stars, hence it should be possible to determine the sky background from the object images. No separate sky observations were obtained by the observers. The number density of objects is higher than for the field of Abell 1689 but because most of the objects are stars instead of galaxies, the fraction of pixels affected by object light is lower at about 1.2%.

The data set consists of 15 exposures obtained during one OB on 2008-09-03 with Hawk-I in the *K* band. The exposures have varying integration times: the first ten exposures have $T_{\mathrm{exp}} = 10\,\mathrm{s}$, the

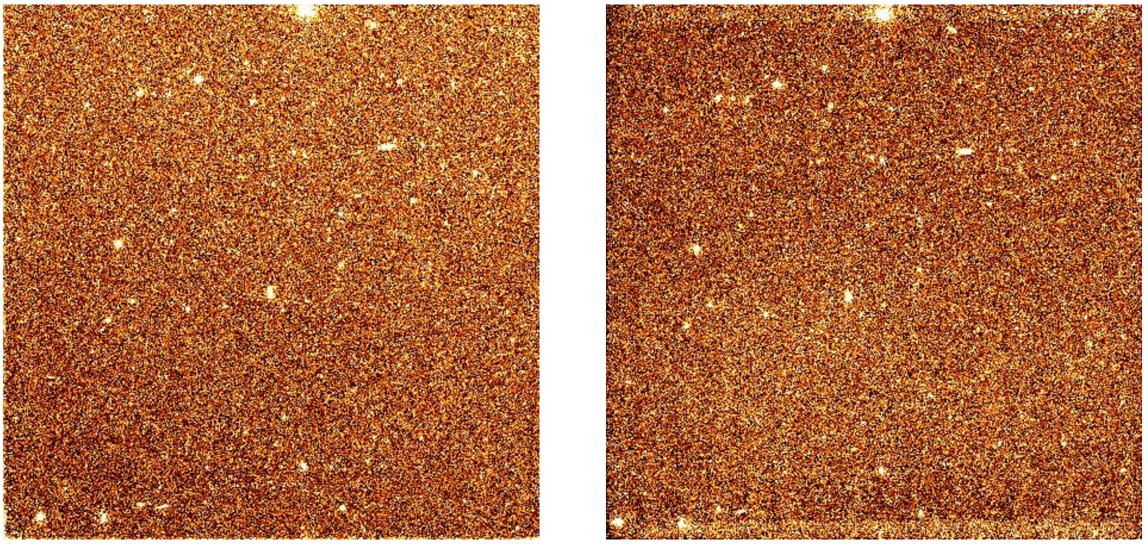

Figure 11: Background corrected images of IC 1613 (HAWK-I, *K* band). Left panel shows exposure 3, right panel exposure 09 of 15. Note the bright stripes at the bottom and top of the latter image. The correction image of this image contains shorter exposures.

remaining five $T_{\mathrm{exp}} = 1.3\,\mathrm{s}$. The data were retrieved from the archive and prereduced using the Hawk-I pipeline.

The result of the subtraction on a single exposure is shown in Fig. 11. The left panel shows an exposure whose correction image is made up solely of 10 s exposures whereas the correction image of the exposure shown in the right panel includes short exposures as well. The background structure of the long and short exposures appear not to match up, which shows up as residuals at the chip edges. This is not the case for exposures where the correction image contains only long exposures.

9.4 HAWK-I: M 4, J band

M 4 is a globular cluster whose core region has a diameter of $\sim 10\,\mathrm{arcmin}$. The cluster is essentially resolved into individual stars in the Hawk-I observations, but the star density is very high, the fraction of pixels affected by star flux exceeding 20 % in pointings including the cluster centre.

Unfortunately, the WCS of the files is corrupted, with wrong `CRVAL` and missing `CDi_j` entries. We managed to use the header files created by Armin Gabasch during recent astrometry tests with `scamp` to correct the WCS for most of the exposures; for some, no header files could be found. The correctness of the WCS was checked by comparison with DSS images.

The available data come from two OBs named “M4_GC” and “M4_GC-B” executed on consecutive nights 2007-08-02/03 and 2007-08-03/04. After WCS correction, there are 23 exposures from the first, and 21 exposures from the second OB.

This data set demonstrates that a too high object density touches the limits of the capabilities of the recipes. The mosaic (Fig. 12) shows that in the central parts of the cluster there are regions where not enough sky information was available for the recipe to work properly, resulting in dark patches in the background. The best option for this target field would have been to obtain separate sky exposures

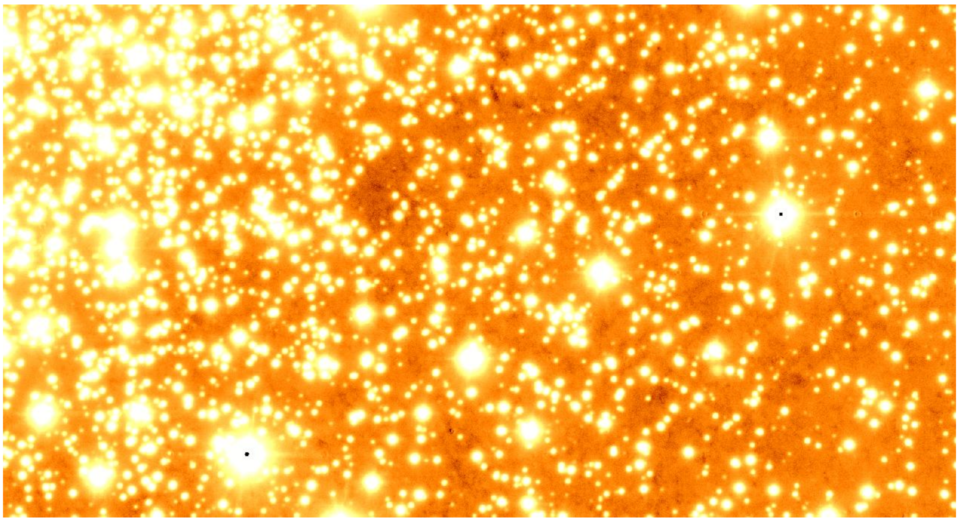

Figure 12: Zoom-in on the background corrected mosaic of globular cluster M 4 (HAWK-I, J band).

as was done for NGC 7793 (Sect. 9.2). On the software level, an improved sky estimate might be obtainable if the robust estimators of Sect. 7 were replaced by a classification method. Even in cases where more than half of the exposures are affected by object flux in a given pixel, such that even the median fails to yield a value that is representative of the background level, one would expect the distribution of values to be bimodal. The lower mode arises from those exposures that have only background flux whereas the higher mode is from exposures affected by object flux. A classification method would identify these modes and assign to each data value a probability that it belongs to one or the other mode. A simple classification method would fit two Gaussians to the sample.

9.5 ISAAC: CDFS F01, K_s band

As an example of ISAAC data, we test the sky subtraction recipe on a set of exposures of the Chandra Deep Field South. We randomly choose OB 216257, observed on 2006-01-07. There are 40 object exposures, taken in the K_s band with integration time 15 s. The field is uncrowded with object fraction below 1 %.

We use these images to check how measurements on objects vary from exposure to exposure through the sequence and how sky correction affects the measurements. SExtractor is run on the 40 exposures, once on the basic calibrated exposures, and once on the sky-corrected exposures (created using the two-pass method with $\kappa\sigma$ clipping with $\kappa = 3$). We choose a photometric parameter (total magnitude `MAG_AUTO`) and a shape parameter (half-light radius `FLUX_RADIUS`) for the comparison.

The upper panel of Fig. 13 shows the magnitude measurements across the 40 basic calibrated exposures. While the measurements from the individual exposures are consistent with each other, there is substantial scatter from exposure to exposure. The lower panel of Fig. 13 shows magnitudes measured on the sky corrected images. The scatter in the lines is now substantially reduced compared to the upper panel. Fig. 14 plots the standard deviation of the measurements (for objects with at least four detections) as a function of the object's mean magnitude.

The situation for the half-light radii is shown in Fig. 15. Again, the scatter in the measurements is significantly reduced by the sky correction, to the extent that a trend of deteriorating seeing is apparent in the half-light radii of stars (the group of lines with the smallest radii) increasing from 2 pixels in the first exposure to 3 pixels in the last. Fig. 14 plots the standard deviation of the measurements (for objects with at least four detections) as a function of the object's mean radius.

9.6 ISAAC: CDFS, H band

Stacks of the full set of ISAAC data on individual tiles are shown in Fig. 17. These stacks are created from ~ 300 exposures each. Due to the vastly decreased noise compared to the individual exposures, coherent structures in the background due to local over- or under-subtraction are much more readily apparent in combined images. A mosaic of two adjacent tiles is shown in the report *DR01 - Astrometry*, demonstrating the homogeneity of the background over the mosaic.

10 Conclusions

We have described a set of recipes and functions based on the Common Pipeline library that estimate and subtract sky and other background contributions from astronomical images from instruments that operate in the optical and near-infrared wavebands.

The near-infrared routines have been developed and tested on a number of data sets from HAWK-I and ISAAC. The optical (smooth) sky correction has been tested on artificial data.

We have also discussed several estimators for the mean of a sequence of noisy pixel values, including robust estimators and outlier rejection algorithms. The use of these estimators and algorithms is not restricted to the stacking of images on the image grid for the purpose of creating a sky estimate but can also be used in stacking of astrometrically corrected images. A simple form is given by the recipe `esosoft_shift_and_add` which is part of the two-pass sky correction. We intend to make use of them in the final stacking of images after the astrometric calibration, taking over the last step of operation of `swarp` as described in the report *DR01 - Astrometry*. Finally, these estimators will form the core of the recipe for stacking one-dimensional spectra (see *DR03 - Stacking of 1D spectra*).

References

Bertin, E. 2008, SWarp v2.17.0, Institut d'Astrophysique & Observatoire de Paris

—. 2009, SExtractor v2.5, User's manual, Institut d'Astrophysique & Observatoire de Paris

Bertin, E. & Arnouts, S. 1996, A&AS, 117, 393

Björck, Å. 1996, Numerical methods for least squares problems (Society for Industrial and Applied Mathematics)

Cuby, J. G., Lidman, C., & Moutou, C. 2000, ESO Messenger, 101, 2

DasGupta, A. 2008, Asymptotic Theory of Statistics and Probability (New York: Springer)

Maritz, J. S. & Jarrett, R. G. 1978, Journ. Am. Stat. Assoc. 73, 194

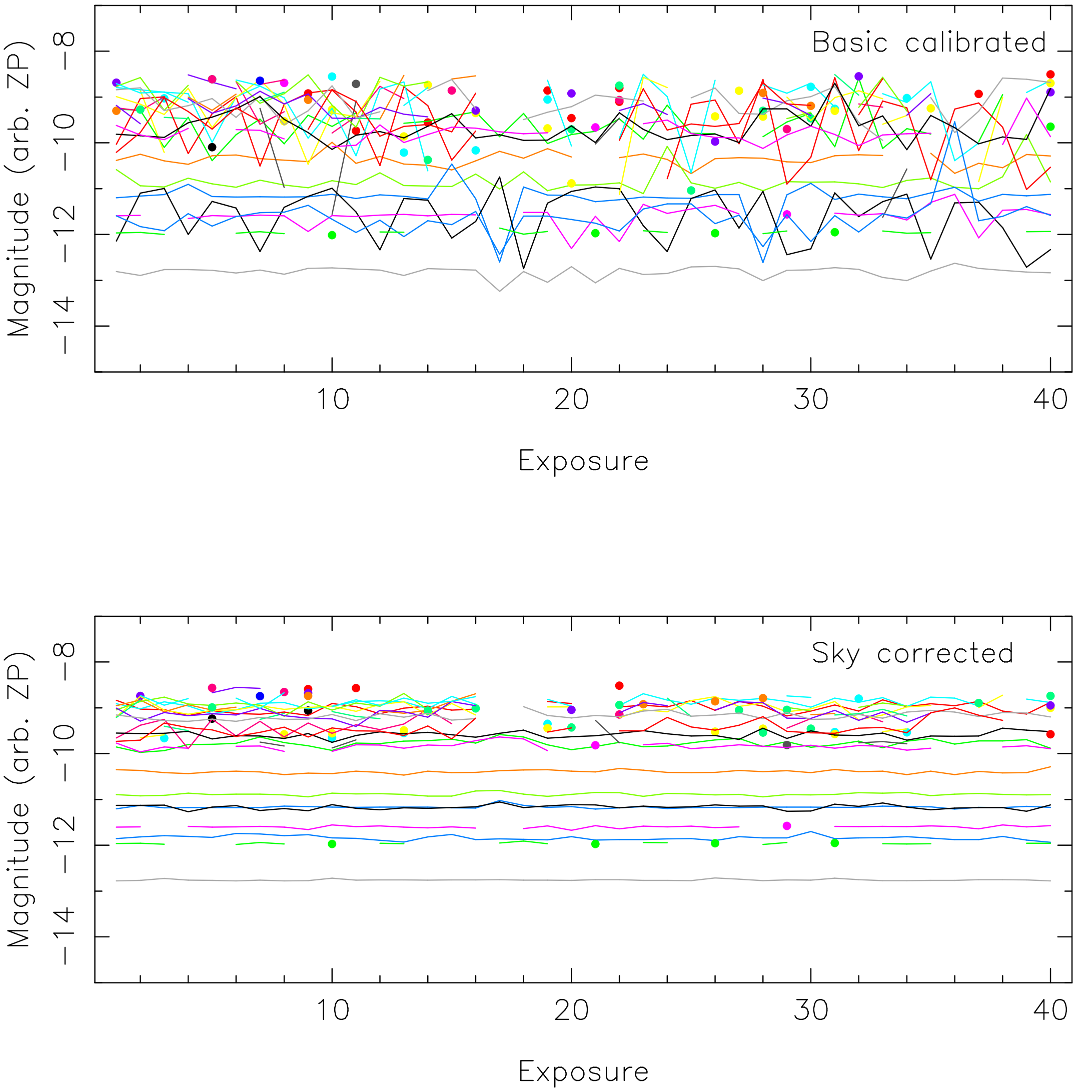

Figure 13: The figure shows how the measured magnitudes vary within the set of 40 ISAAC K_s exposures on the CDFS, field 01, OB 216257. The top panel shows the measurements on the basic calibrated images, before sky subtraction. The bottom panel shows the measurements after sky subtraction. Lines are broken when the corresponding objects are not detected on one or several exposures. Magnitudes are measured with an arbitrary zero point.

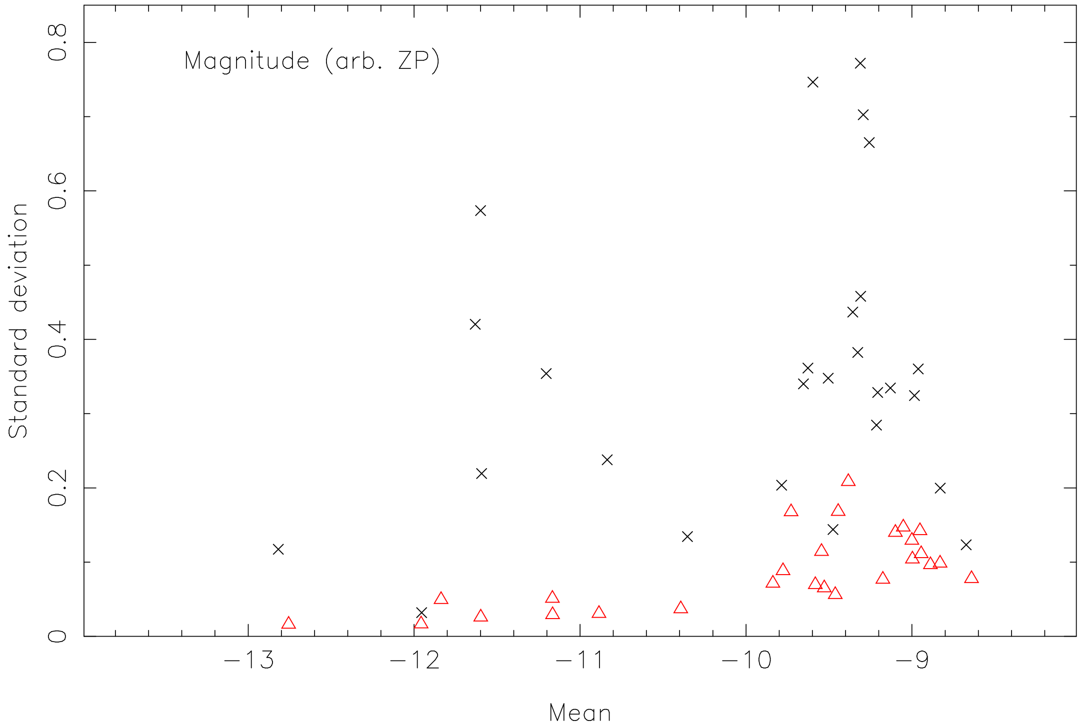

Figure 14: Scatter of magnitude measurements across 40 ISAAC K_s exposures. The crosses are for measurements on the basic calibrated exposures, before sky subtraction, the red triangles for measurements on sky-corrected exposures.

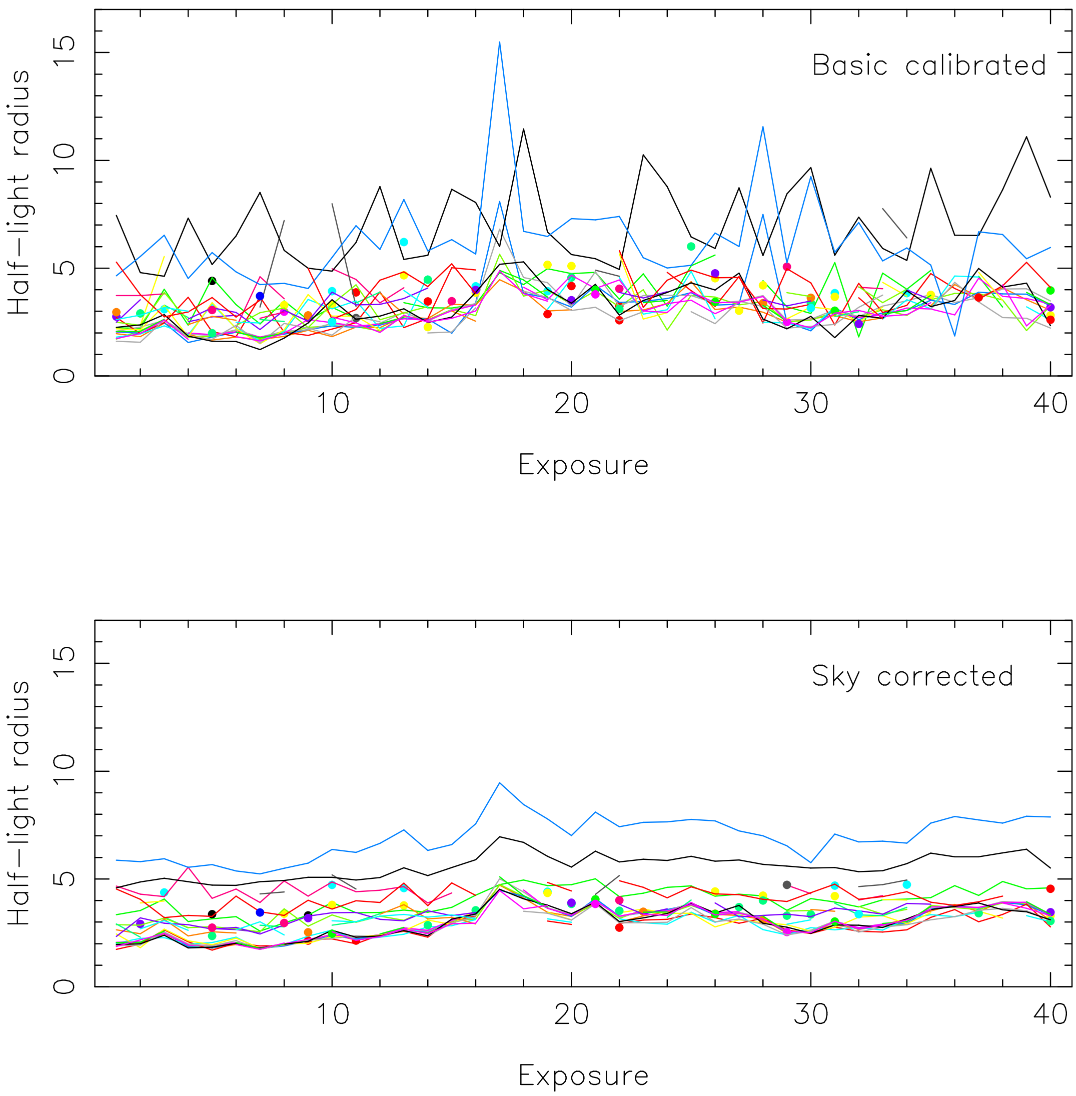

Figure 15: The figure shows how the measured half-light radii vary within the set of 40 ISAAC K_s exposures on the CDFS, field 01, OB 216257. The top panel shows the measurements on the basic calibrated images, before sky subtraction. The bottom panel shows the measurements after sky subtraction. Lines are broken when the corresponding objects are not detected on one or several exposures. The radii are given in units of pixels (0.15 arcsec).

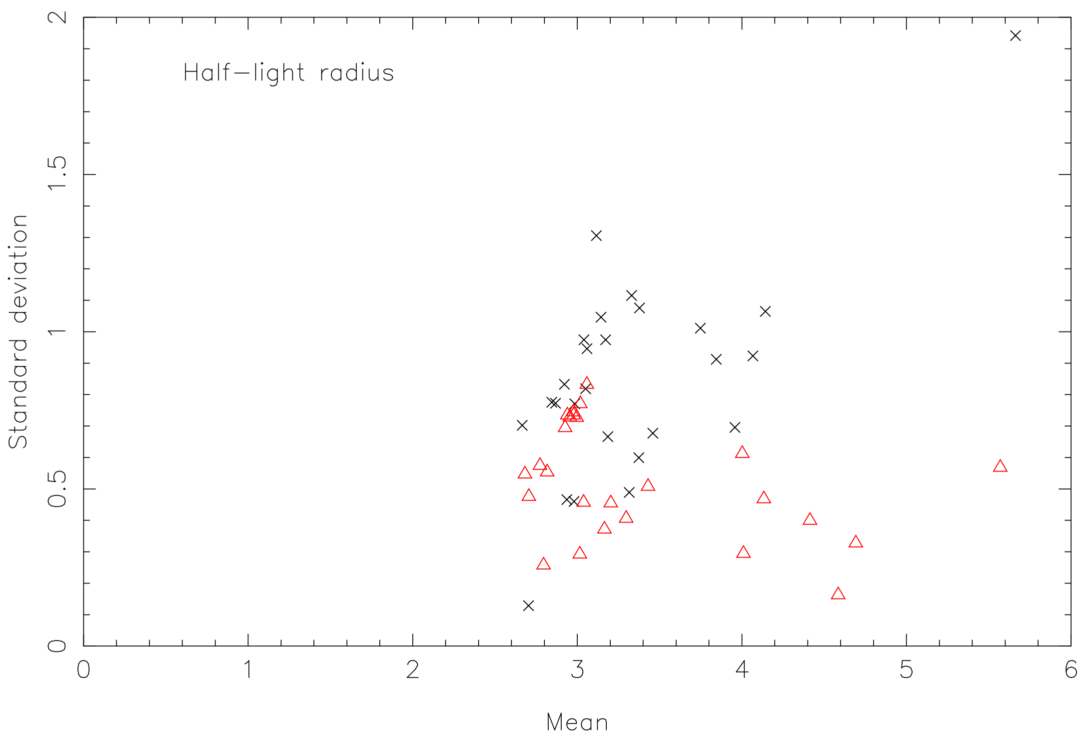

Figure 16: Scatter of measurements half-light radius across 40 ISAAC K_s exposures. The crosses are for measurements on the basic calibrated exposures, before sky subtraction, the red triangles for measurements on sky-corrected exposures.

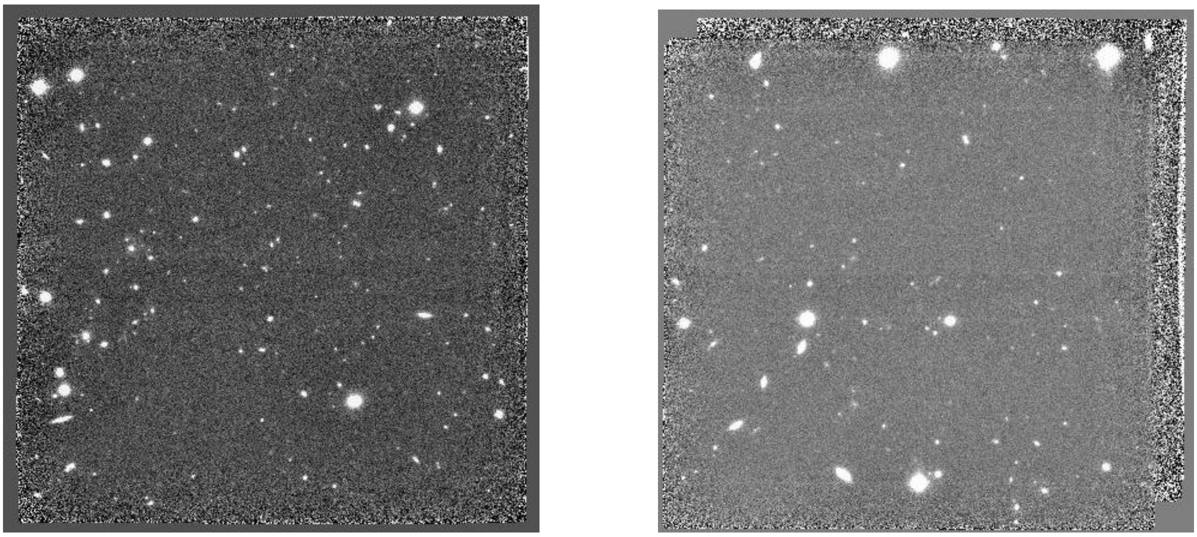

Figure 17: ISAAC H-band stacks of tiles F06 (left) and F18 (right) of the GOODS-South field. Each image is created from ~ 300 individual exposures.

Marmo, C. 2007, in Astronomical Society of the Pacific Conference Series, Vol. 376, Astronomical Data Analysis Software and Systems XVI, ed. R. A. Shaw, F. Hill, & D. J. Bell, 285

Newberry, M. V. 1999, in Astronomical Society of the Pacific Conference Series, Vol. 189, Precision CCD Photometry, ed. E. R. Craine, D. L. Crawford, & R. A. Tucker, 74–+

Osterbrock, D. E., Fulbright, J. P., Martel, A. R., Keane, M. J., & Trager, S. C. 1996, Publ. Astr. Soc. Pac. , 108, 277

Patat, F. 2004, ESO Messenger, 115, 18

Rider, P. R. 1960, Journ. Am. Stat. Assoc. 55, 148

Vandame, B. 2004, PhD thesis, Université de Nice – Sophia Antipolis