

Stellar Photometry and Astrometry in Adaptive Optics Images

Emiliano Diolaiti

Dipartimento di Astronomia, Università di Padova, Vicolo dell'Osservatorio 5, I-35122 Padova, Italy

Dipartimento di Astronomia, Università di Bologna, Via Ranzani 1, I-40127 Bologna, Italy, Email: diolaiti@bo.astro.it

Douglas G. Currie

European Southern Observatory, Karl Schwarzschild-Str. 2, D-85748 Garching b. München, Germany

Abstract. Adaptive optics techniques allow high resolution imaging of crowded stellar fields. The photometric and astrometric information must be extracted from the data with suitable post-processing methods. From this point of view adaptive optics images involve new problems: the StarFinder code, described in this lecture, represents an attempted solution.

1. Introduction

A stellar field can be thought of as a superposition of stellar images on a smooth background emission (see Fig. 1). Each star, as seen through the turbulent atmosphere partially corrected by an Adaptive Optics (AO) system, appears as a copy of an intensity pattern (Point Spread Function or PSF) resulting from the residual atmospheric distortion and the telescope diffraction. In general the PSF changes across the field of view, so that two stars in different directions in the sky will be imaged more or less differently. We will discuss briefly this problem, referred to as angular anisoplanatism, in the last section of the lecture. For most of our considerations we will assume the PSF to be reasonably uniform, a condition which is normally met in practice.

The image analysis involves the detection of the stellar sources and the determination of their position (astrometry) and relative magnitude (photometry). The main difficulties arise from the overlapping of the PSFs of neighboring stars, which contaminate each other (see Fig. 2). In the case of AO images the problem is still more difficult to solve, because of the complicated structure of the PSF, which generally shows a sharp peak, one or more fragmented diffraction rings and an extended irregular halo. These features may lead to false detections, loss of faint stars and worsening of the photometric accuracy. The troubles and pitfalls related to the analysis of images with highly structured PSF have been described, for instance, by Esslinger & Edmunds (1998) for AO observations and Sosin & King (1997) in the specific case of HST images taken with the Faint

Object Camera, an instrument which highlights the features of the HST PSF due to its small pixel size.

The next section describes how these problems have been faced in the development of the StarFinder code (Diolaiti et al. 2000).

2. General description of the algorithm

2.1. Guidelines

From the introductory discussion we can derive some useful hints as to handle the problems related to the overlapping of the stellar images and the complicated shape of the PSF.

The first guideline is to analyze the presumed stars in order of decreasing intensity, according to the principle that a bright star can be examined neglecting the contamination due to the fainter sources around, at least to a first approximation. If we keep track of what we are doing, when we examine a faint object we may take advantage of the already available knowledge on the formerly analyzed brighter sources. In this way we reduce the probability of taking spurious PSF features for true stars. StarFinder keeps track of the bright sources by means of a *synthetic field* generated by placing a suitably scaled copy of the PSF for each detected star. This fake image is continuously updated and becomes more and more similar to the observed frame as the analysis proceeds.

The second guideline is related to the technique used to determine the photometry and astrometry of the stars. A commonly used method is the so called multi-component PSF-fitting, which approximates a group of stars with a parametric model, representing each source as a copy of the PSF; the free parameters to be optimized are the positions and fluxes of the stars. A question may arise about the number of sources to be fitted together. In principle there is no upper bound: in a very crowded field there may be groups of several tens overlapping sources, contaminating each other, and it may be inefficient to fit them all together. A possible strategy is to perform a single-component fit of each source, provided the available information on the formerly analyzed stars is taken into proper account. The necessary information on the brighter sources may be extracted again from the synthetic field. A final re-fitting improves the overall accuracy, since also the bright sources are more or less contaminated by the fainter ones, initially neglected. Of course there are exceptions to the single-component fitting strategy: very compact groups at separations comparable to the PSF FWHM should be examined with a multi-component model.

The last guideline is related to the PSF approximation to be used. The accurate knowledge of its features outside the central peak (see Fig. 2) is fundamental to perform a deep study of the stellar field, achieving accurate photometry of faint stars and avoiding either false detections or star loss. A seeing-limited PSF can be represented by parametric models. The HST PSF has been often expressed as the sum of an analytic model, generally used to drive the fitting of the star positions, and a residual map, to account for the actual shape. The approach implemented in StarFinder is to use a digital template of the PSF with no analytic modeling. This allows a very accurate representation of all the features that might cause the problems described above. It is intuitive that, in order to match a given star, the PSF array must be shifted by sub-pixel offsets,




Figure 1. Left: simulated stellar field with 1000 stars, represented by scaled copies of an experimental PSF placed at random positions. Notice the features in the halo of the brighter stars, which may lead to false detections. Right: same field contaminated by a smooth non-uniform background nebulosity.



Figure 2. Left: experimental PSF in K band; the Strehl ratio is about 40%. Notice the diffraction pattern and the irregular structures in the halo. The PSF has been oversampled for a visualization purpose: the FWHM of the peak in the original frame is ~ 3.6 pixels. Right: simulated binary star given by the superposition of two copies of the PSF with a flux ratio of 20 and a separation of ~ 0.5 arcsec. It is evident that in order to analyze the fainter component it is necessary to take into account the halo contamination of the brighter one.

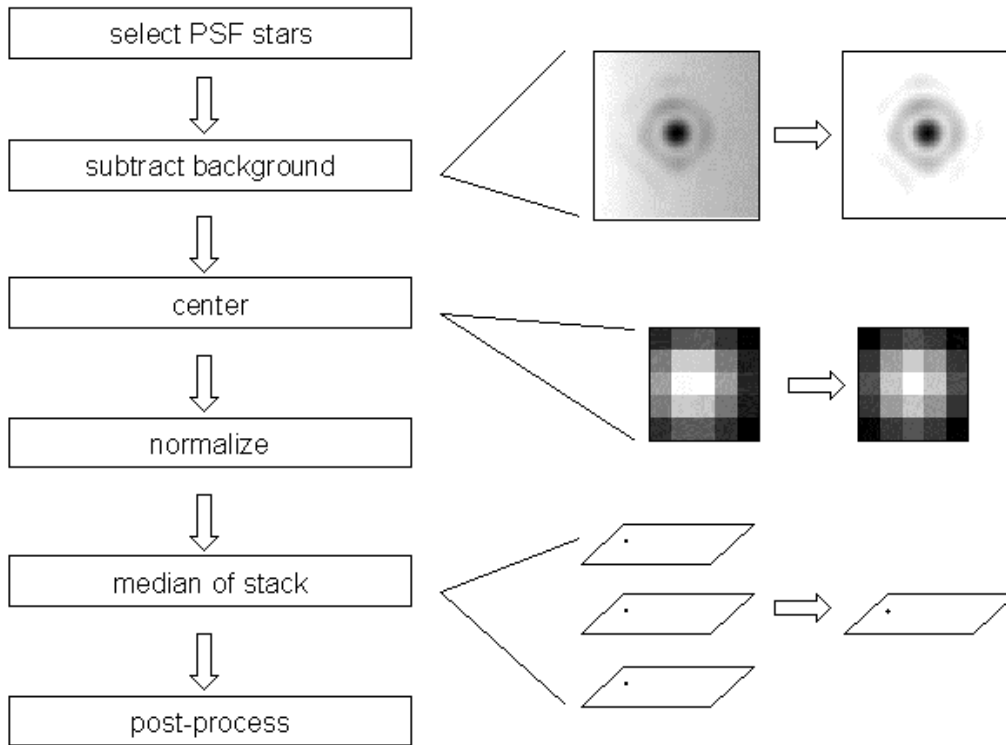


Figure 3. Left: PSF extraction flow-chart. The selected stars may also be decontaminated from the secondary sources around: this optional step is not shown here. Right: some operations are illustrated in more detail.

an operation that can be accomplished by means of interpolating techniques in the case of AO images, which generally fulfill the critical sampling condition (see Sect. 3.5.).

2.2. PSF extraction

The accuracy of the PSF estimate is essential in StarFinder, because the PSF array is used as a template for all the stars in the image to be analyzed. A reliable way to estimate it is to combine the images of some stars in the observed field (see Fig. 3). The user selects the most suitable ones, which are background subtracted, possibly cleaned from the contaminating sources around, centered with sub-pixel accuracy, normalized and superposed with a median operation. The centering is performed by an iterative shift of the stellar image in order to cancel the sub-pixel offset of its centroid: this step, involving the interpolation of the array, can be reliably performed only on well-sampled data (see Sect. 3.5.). The median operation, which is performed pixel-by-pixel, is preferred to the mean because it is less sensitive to anomalous pixels (outliers) which might appear in one or more stellar images among the selected ones. The retrieved PSF may be post-processed in order to reject any spurious residual features and smooth the noise in the extended halo.

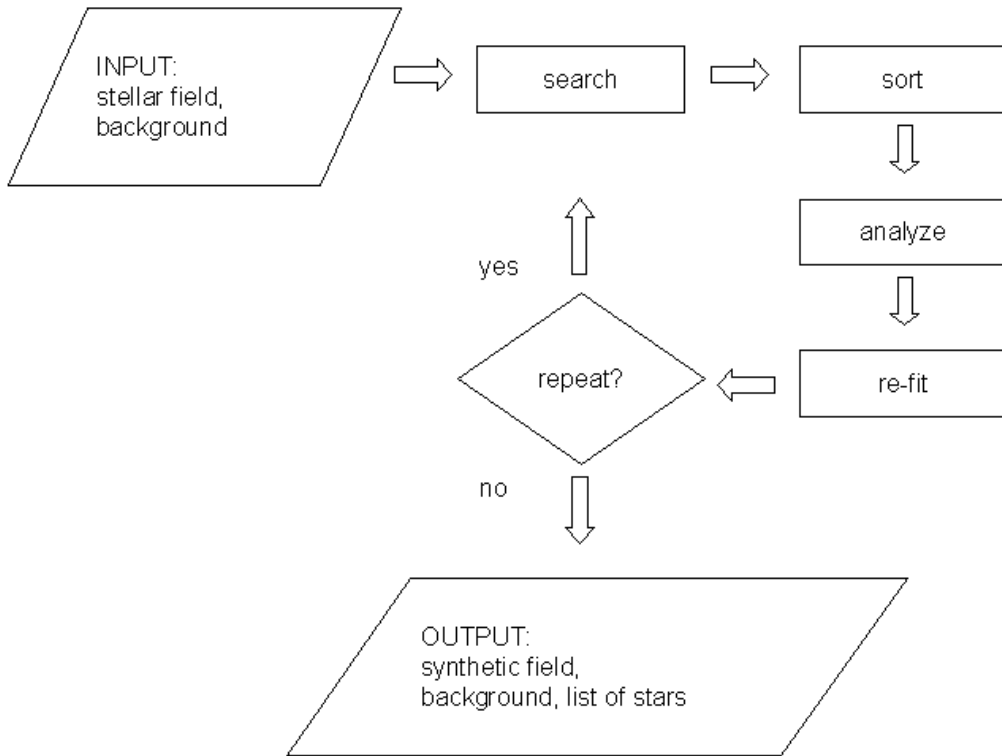


Figure 4. Main loop flow-chart. The presumed stars in the image are sorted by decreasing intensity and analyzed one by one; for each of them the available information on the already detected brighter stars is taken into account. The last step is a re-fitting of the successfully detected stars out of the original list. Then the algorithm may be stopped or a new iteration of the loop may be started, searching for possible lost stars in the residual image.

In a very crowded field the PSF estimate is more reliable if the selected stars are decontaminated from the secondary sources around. This task can be accomplished manually, selecting the stars to subtract, or automatically. In the latter case the PSF estimate is used to perform a preliminary image analysis with the ultimate goal to detect and subtract the contaminating sources; then the PSF extraction procedure is applied again to the decontaminated PSF stars. This scheme may be applied iteratively.

The PSF extraction procedure reconstructs the core of the saturated stars by replacing the corrupted pixels with the central part of a preliminary PSF estimate. Accurate positioning is achieved by means of a cross-correlation technique, while the scaling factor is determined with a least squares fit to the wings of the star to repair. For a detailed description see Sect. 3.2..

If the PSF template cannot be extracted directly from the field, due to extreme crowding or lack of suitable stars, StarFinder can still be applied using a PSF estimated by means of other methods, as the reconstruction technique proposed by Véran et al. (1997).

2.3. Detection of the stars, photometry and astrometry

A flow-chart illustrating the sequence of operations for the photometric and astrometric analysis of the stellar field is shown in Fig. 4.

At first we build a list of objects, the presumed stars, which satisfy the detection condition

$$i(x, y) > b(x, y) + t \quad (1)$$

where $i(x, y)$ is the observed intensity, $b(x, y)$ the background emission and t a fair detection threshold, related to the noise standard deviation.

The presumed stars are analyzed one by one in order of decreasing intensity. To illustrate a generic step of the algorithm, we consider the $(n + 1)$ -th object in the list, after the examination of the first n . A small sub-image of fixed size is extracted around the object (Fig. 5, left). This patch may contain brighter stars formerly analyzed, fainter objects neglected in the current step and features of other stars lying outside. The information on the brighter sources is recorded in the synthetic stellar field (Fig. 5, center), defined as the sum of two terms: a superposition of PSF replicas, one for each star detected up to this point, and an estimate of the background, assumed to be non uniform in general. The synthetic field contribution is subtracted from the sub-image and if a statistically significant residual is found (Fig. 5, right) it is compared to the central core of the PSF by means of a correlation check. If the correlation coefficient is higher than a pre-fixed threshold, the object of interest is rated similar to the PSF and accepted. The accurate determination of its position and relative flux is attained by means of a local fit, in which the observed sub-image is approximated with the multi-component model described in Sect. 3.5.. If the fit is acceptable, the parameters of the new detected star are saved and those of the already known sources, which have been possibly re-fitted, are upgraded. The new star and an upgrade of the re-fitted sources are added to the synthetic field.

When all the objects in the initial list have been examined, the successfully detected stars are fitted again, this time considering all the known sources, not only the brighter ones. This step ensures better photometric and astrometric accuracy; it may be iterated a pre-fixed number of times or until a convergence condition is met.

At the end of the analysis it is possible to stop the algorithm or search for possibly lost object by removing the detected stars and an upgraded background estimate. It should be stressed that this image subtraction is just a tool to highlight significant residuals: any further operation is performed on the original frame, in order to take into account the effects arising from the superposition of the PSFs of neighboring sources. Generally, after 2-3 iterations of the main loop, the number of detected stars approaches a stable value.

2.4. Blends

We consider the simplest example of a blend: a binary star. If the two components of the binary are well separated and their intensity peaks are clearly discernible, the code analyzes them with the standard procedure described in the previous sub-section. Otherwise the secondary component is not detected as a relative maximum and it is lost (see example in Fig. 6).




Figure 5. Analysis of an object (see also the *analyze* step in Fig. 4). Left panel: the sub-image centered on the object of interest (1) includes an already known brighter star (2), a fainter object (3) neglected in the current step and the PSF features of bright stars outside, represented by the structures in the upper part. Central panel: the available information on the stars detected so far is extracted from the synthetic field. Right panel: the subtraction of the synthetic field (central panel) from the image (left panel) suggests that the central peak (1) is not a spurious PSF feature of a bright star, so it is worth fitting it.

If the separation is not too small, a further iteration of the main loop generally enables the algorithm to detect the fainter component by subtracting the brighter one (Fig. 6, upper row). This strategy forces the two stars to pass the correlation test, the principal component as a presumed single object and the secondary in a subsequent iteration. It may fail when the two components are almost equally bright and have a separation close to $1/2$ FWHM, which represents a kind of lower limit: the residual corresponding to the secondary after subtracting the principal component may have a distorted shape (Fig. 6, lower row) and might not pass the correlation check.

The latter case may be handled by a method based on a thresholding technique. The object is cut at a prefixed level, about 20% below the central peak, and transformed to a binary array, setting to 1 all the pixels above the threshold and to 0 the pixels below. If the area of the pixels with value 1 is more extended than the PSF, the object is considered a blend and the secondary star may be detected by subtracting the brighter one; then a two-component fit allows accurate astrometry and photometry of the two sources. This algorithm can be iteratively applied to multiple blends. It should be stressed that the area measurement is not reliable when the value of the cutting threshold is comparable to the noise level: for this reason the de-blending procedure is applied only to objects with a suitable signal-to-noise ratio. Moreover the measurement is reliable when the data are adequately sampled. The de-blending procedure is applied at the end of the last iteration of the main loop, when all the resolved sources have been detected: in this way the probability that a single object may appear artificially blurred because of the spurious contamination of still unknown sources is largely reduced.

In a normal case, like the simulated field of Sect. 4.1., two or three iterations of the main loop find almost all ($\sim 99\%$) the stars that StarFinder may detect. The




Figure 6. Examples of binaries. The upper row refers to a separation of 1 FWHM and a flux ratio of 2:1; in the lower row the separation is 0.75 FWHM and the flux ratio 1:1. The binaries are first fitted with a single component, which is then removed highlighting the residual corresponding to the fainter star. When the binary is resolved according to the Rayleigh criterion (upper row) the residual has a nice shape, very similar to the PSF. When the binary is resolved according to the Sparrow criterion instead (lower row) the residual is very distorted and an attempt to match it to the PSF will not be successful.

de-blending procedure described above adds $\sim 1\%$ more stars, without additional false detections. Normally we perform two or three iterations of the main loop and apply the de-blending strategy only in very crowded fields.

3. Details

3.1. Background estimation

A reliable estimate of the background is necessary to define the detection condition (Eq. 1) and to compute the correlation coefficient of the presumed stars with the PSF.

A straightforward technique to estimate the background is to smooth the image by median filtering, replacing each pixel with the median computed over a suitable neighborhood, of size comparable to the characteristic width of a stellar image. This method tends to over-estimate the background underneath strong peaks. A more accurate technique is described in Bertin & Arnouts (1996) and




Figure 7. Background estimation. a) stellar field; b) partitioning; c) array of local background measurements; d) smoothing; e) interpolation onto the same grid of the original image; f) background-removed field.

illustrated in Fig. 7. The image is partitioned into a grid of sub-regions and a local estimate is calculated for each patch by means of an IDL implementation of the DAOPHOT SKY routine (Stetson 1987), due to Landsman (1995). This array of sky measurements is smoothed by median filtering and interpolated onto the same grid of the input image.

It should be stressed that the background computation is unavoidably affected by the presence of bright sources. In general a more accurate estimate can be obtained after the analysis of the stellar field, when the most contaminating sources are known and can be temporarily subtracted.

3.2. Saturated stars

Saturated stars provide precious information on the PSF halo, so it may be useful to include them, appropriately reconstructed, in the PSF extraction process. In addition the repaired saturated stars can be recognized by the star detection algorithm and their contribution be taken into account when analyzing fainter sources.

The core of a saturated star is replaced with a shifted scaled replica of a preliminary estimate of the PSF. The repaired star is defined as

$$i_{rep.}(x, y) = \begin{cases} i(x, y) & \text{if } i(x, y) < T \\ fp(x - x_0, y - y_0) & \text{otherwise} \end{cases} \quad (2)$$

where T is the upper linearity threshold of the detector. The position of the center (x_0, y_0) is estimated by maximizing the correlation coefficient, which is not sensitive to the intensity levels. The scaling factor f is calculated with a least squares fit to the wings of the saturated star. Of course the saturated pixels are excluded from all the computations. The background is temporarily subtracted in order to prevent affecting the repair process.

This procedure has been applied to the brightest star in the Galactic Center image shown in Sect. 4.2.. This source is not saturated in the original data, but it has been artificially corrupted by an upper cut at half maximum. The repair procedure is able to reconstruct the original peak with an error of $\sim 1.5\%$, imposing a suitable positioning accuracy for the correlation maximization.

3.3. Noise estimation

The estimate of the noise is useful to define the detection threshold (Eq. 1) and to compute the formal errors on the retrieved astrometry and photometry. In general the noise standard deviation is given by

$$\sigma = \sqrt{\sigma_{\text{photon}}^2 + \sigma_{\text{instr}}^2} \quad (3)$$

where the photon noise standard deviation σ_{photon} is proportional to the square root of the number of photons and σ_{instr} is the instrumental noise standard deviation. The overall effect of photon and instrumental noise can be computed if the required parameters (detector gain, read-out-noise, dark current, etc.) are known. Otherwise an estimate of the noise level can still be obtained by means of histogram fitting techniques (Almoznino, Loinger, & Brosch (1993), Bijaoui 1980). Assuming that the intensity of the sky radiation is distributed normally around a typical value, the histogram of the observed intensity levels should be quite similar to a gaussian distribution, whose mode and standard deviation represent respectively the sky value and the associated noise, resulting from the photon signal and the instrumental effects. Actually the contamination due to the stellar sources produces a high-intensity tail and an artificial broadening of the histogram (see Fig. 8, left). This problem can be partially overcome by removing the signal from the image, leaving only the pixel-to-pixel variations associated to pure noise: a reasonable estimate of the signal to subtract for this purpose can be obtained by filtering the data with a median technique (the same described in Sect. 3.1.), using a small smoothing box. After this operation, the histogram is symmetric around its mode and the background noise standard deviation can be estimated by means of a gaussian fit (see Fig. 8, right).

3.4. Correlation coefficient

False detections, associated to noise spikes or residual PSF features of bright stars, are recognized on the basis of their low correlation coefficient with the PSF, which represents a template for each true star in the image. The correlation coefficient (Gonzalez & Woods 1992) is defined as

$$c(a, b) = \frac{\sum_{x,y} [i(x, y) - \bar{i}] [p(x - a, y - b) - \bar{p}]}{\sqrt{\sum_{x,y} [i(x, y) - \bar{i}]^2} \sqrt{\sum_{x,y} [p(x - a, y - b) - \bar{p}]^2}} \quad (4)$$




Figure 8. Left: histogram of the intensity values of a stellar field. The histogram mode is related to the mean background level; the high intensity tail is due to the stellar sources. Right: histogram of the same image after removing the signal, leaving only the pixel-to-pixel fluctuations associated to the 0-mean background noise. Notice the different intensity ranges in the two plots.

where $i(x, y)$ and $p(x, y)$ are the object and the PSF respectively, \bar{i} and \bar{p} are the corresponding mean values. Maximizing the correlation coefficient as a function of the offset (a, b) yields an objective measure of similarity. After maximizing $c(a, b)$ for integral offsets, it is possible to repeat the procedure for sub-pixel shifts, improving the positioning accuracy.

The correlation coefficient is computed on the core of the PSF: typically the central spike of the diffraction pattern is considered, out to the first dark ring. A fair correlation threshold must be fixed in order to discriminate and reject unlikely detections, without losing faint stars contaminated by the background noise: a value of 0.7 or 0.8 is acceptable in most cases.

The correlation coefficient represents also an effective tool for the *a posteriori* selection of the stars with the highest photometric reliability, since generally a very high correlation value is associated to resolved single sources.

3.5. Photometry and astrometry

A sub-image centered on the star of interest is extracted and approximated with the model

$$m(x, y) = s_0(x, y) + \sum_{n=1}^{N_s} f_n p(x - x_n, y - y_n) + b_0 + b_x x + b_y y \quad (5)$$

where $s_0(x, y)$ is the fixed contribution of the known stars outside the sub-image support, extracted from the synthetic field, N_s is the number of point sources within the sub-image, x_n, y_n, f_n are the position and flux of the n -th source, $p(x, y)$ is the PSF and b_0, b_x, b_y are the coefficients of a slanting plane representing the local background. In general the linear approximation of the background adopted here is more accurate for this purpose than the estimate described in Sect. 3.1.. The actual size of the fitting region is comparable to the diameter of the first bright diffraction ring of the PSF; this choice ensures that

the unique information represented by the shape of a high-Strehl PSF is included in the fitting process to achieve better accuracy and prevents the growth of the number of sources to be fitted together. In practice a multi-component fit is performed when the star is in a very compact group, with typical separations comparable to the PSF FWHM. Otherwise a single-component fit is enough: the contamination of the other sources is accounted for by adding the fixed contribution derived from the synthetic field and by the final re-fitting (Sect. 2.3. and Fig. 4).

It should be stressed that the retrieved astrometry and photometry are referred to the absolute centering and normalization of the PSF array. If the PSF has an offset of 1/4 pixel, for instance, all the estimated positions will have the same offset in the opposite direction. If the PSF estimate is truncated or has some spurious features, then the estimated fluxes will be under- or over-estimated by a proportionality factor, which will translate into a zero-point offset of the magnitude scale.

The optimization of the free parameters in Eq. 5 is performed by minimizing the least squares error between the data and the model, by means of an iterative Newton-Gauss technique (Beck & Arnold 1977, Luenberger 1984, Bendinelli et al. 1987, Bevington & Robinson 1994), which is well suited to nonlinear fitting. In general a Newton-like method would require the first and second order derivatives of the model with respect to the parameters. Assuming that the starting guess of the parameters to be estimated is sufficiently close to the solution, a variant of the method can be applied, requiring only the first-order derivatives. Most of the parameters in Eq. 5 yield linear dependence; the only exception are the star positions. If the PSF were represented by an analytic model, e.g. a gaussian in the simplest case, it would be possible to compute the derivatives of the model with respect to the positions with an exact formula. In this case the PSF is represented by a digital template; however it can be translated to the required position applying the discrete version of the Fourier shift theorem

$$p(x - x_n, y - y_n) = FT^{-1} \left[FT [p(x, y)] e^{-i2\pi(ux_n + vy_n)/N} \right] \quad (6)$$

where FT represents the discrete Fourier transform operation, N is the sub-image size and u, v are spatial frequencies. The derivative with respect to x_n is then

$$\begin{aligned} \frac{\partial p(x - x_n, y - y_n)}{\partial x_n} &= \\ FT^{-1} \left[FT [p(x, y)] e^{-i2\pi(ux_n + vy_n)/N} \left(-i \frac{2\pi u}{N} \right) \right] &= \\ FT^{-1} \left[\left(-i \frac{2\pi u}{N} \right) FT [p(x - x_n, y - y_n)] \right] & \quad (7) \end{aligned}$$

and requires in practice an interpolation of the PSF to compute $p(x - x_n, y - y_n)$. A similar technique has been described by V eran & Rigaut (1998). In principle this interpolation-based method can only be applied to Nyquist-sampled (loosely speaking: well-sampled) images and this is currently the main limit of

the algorithm. We recall here that a band-limited function in one dimension is Nyquist-sampled if the step size fulfills the condition

$$\Delta x \leq \frac{1}{2}u_c \quad (8)$$

where u_c is the so-called cut-off frequency of the spectrum. In this case it is possible to reconstruct the original continuous function from a set of equally-spaced samples by means of the so-called sinc interpolation. For a nearly diffraction-limited PSF, which is band-limited in two dimensions, the critical sampling condition is generally stated by saying that its FWHM must measure at least two pixels. Many accurate and efficient interpolation schemes exist to perform the PSF shift required by the fitting procedure. StarFinder applies the cubic convolution interpolation method (Park & Schowengerdt 1983), which approaches very closely the optimal sinc interpolation for well-sampled data; this algorithm is implemented in the IDL function INTERPOLATE.

Even though the interpolation of the PSF array is allowed only on well-sampled data, the cubic convolution technique seems to produce acceptable results even with marginally under-sampled images. Tests performed on Airy diffraction patterns, with a sampling step twice as large as the critical sampling step size, indicate that the interpolation-induced oscillations amount to a few percent of the image peak, as opposed to 10-20% of other interpolation techniques like Fourier shift or bicubic splines.

3.6. Code structure and interface

The StarFinder code has been provided with a collection of auxiliary routines for data visualization and basic image processing, in order to allow the user to analyze a stellar field, produce an output list of objects and compare different lists, e.g. referred to different observations of the same target. A widget-based graphical user interface (Fig. 9) has been created, to make the usage easier. The code is entirely written in IDL language and has been tested on Windows and Unix platforms supporting IDL v. 5.0 or later. It is organized in three levels of abstraction:

- a low level including the pure computational code, i.e. a collection of IDL procedures and functions with normal input/output parameters, but no file input/output and no graphic elements;
- an intermediate level, including many specialized widget-based applications, each associated to a particular task (e.g. PSF extraction, noise estimation, etc.); each of these applications is just a graphical interface wrapping one or more routines of the lower level;
- a high level, represented by the main interface, which handles the global data required by the program and allows the user to call the various tasks forming the intermediate level.

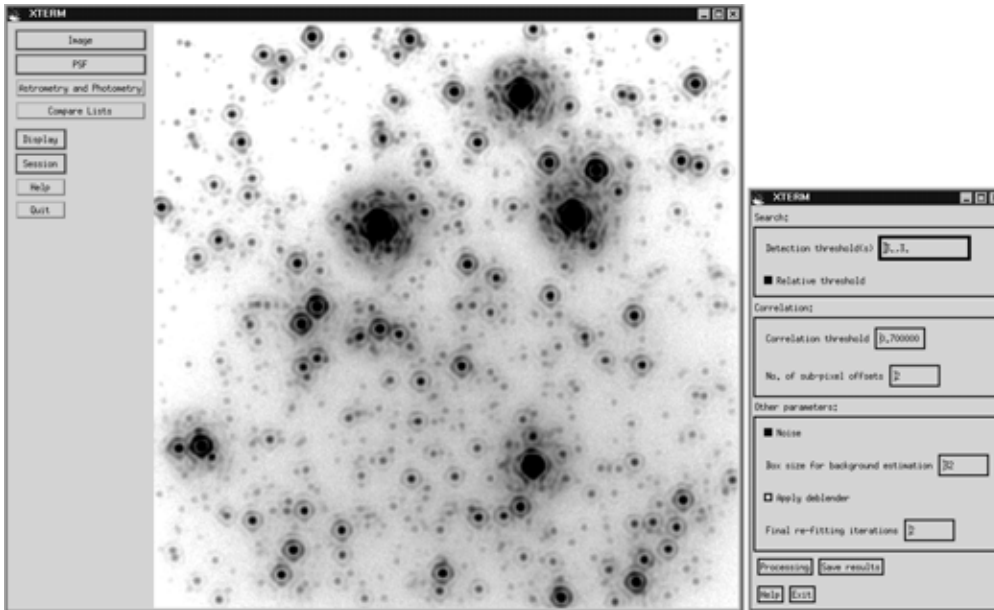


Figure 9. Top-level (left) and intermediate-level (right) user interfaces. The computational code is wrapped by these interfaces.

4. Applications to isoplanatic fields

4.1. Simulated data

The first application is related to a simulated image: in this case the answer is known, so it is an interesting test to compare the list of retrieved objects with the original one, to evaluate the detection reliability and the photometric and astrometric accuracy of the program. The image includes 1000 stars, placed randomly in the frame and with a given magnitude distribution over a range of 8 *mag*. Each star is represented by a scaled copy of a long exposure PSF (see Fig. 2). The image is 368×368 pixels large (13×13 *arcsec*²) and has a stellar density of 6 *stars* \times *arcsec*⁻². Photon, readout noise and a background nebulosity, normalized to the same flux of the stellar sources, have been added to the image. The faintest stars have a peak signal-to-noise ratio of ~ 5 .

With a standard reduction, using the default parameters of the method, it is possible to find almost 90% of the stars. The blending effects and the contamination due to the halo of the four brighter stars account for the lost sources. The false detection cases are negligible (1 out of 1000 stars). The photometric and astrometric accuracy for the detected sources are shown in Fig. 10. About 80% of the detected stars have both astrometric error smaller than 0.1 PSF FWHM and photometric error smaller than 0.1 magnitudes. The stars with less accurate measurements are generally faint or in very compact groups, where the loss of some sources introduces larger errors on the detected ones.

After discussing the performance of the code with a standard analysis, it is interesting to examine how the results are affected by the main parameters of the method. Applying the de-blending strategy, it is possible to detect about




Figure 10. Photometric and astrometric accuracy in the simulated stellar field. Left: estimated vs. true magnitudes of successfully detected stars; the magnitude zero-point on the horizontal axis is arbitrary. The photometry is essentially linear; the points corresponding to negative errors (estimated magnitude brighter than the truth) towards the faint end of the sequence may be attributed to crowding effects: the loss of some stars in blended groups introduces an error in the photometry of the detected ones. Right: astrometric errors, represented by the distance between the calculated and the true position, expressed in FWHM units.

1% more stars. Decreasing the detection threshold from 3 to 1 times the noise standard deviation allows the program to detect some more binaries, but with a higher probability of false detections (10 cases instead of 1, though with $mag > 8$). Increasing the threshold on the correlation coefficient, from 0.7 to 0.8, eliminates false detections, but the number of lost stars increases by about 60%; the additional lost sources are generally fainter than magnitude 7, but not necessarily in crowded groups. With a lower correlation threshold (0.6), it is possible to detect more faint isolated stars and binaries, at separations between 1 and 2 PSF FWHM, but with a slightly higher probability of false detections (2 instead of 1). Finally the astrometric and photometric accuracy approaches a stable level after a few (~ 2) re-fitting iterations.

4.2. Real data: experiment with synthetic stars

The second application is related to a real image: a K-band PUEO observation of the Galactic Center (Fig. 11, left), kindly provided by François Rigaut (Rigaut et al. 1998). The Strehl ratio in this image is $\sim 45\%$ and the pixel size is about half the critical sampling step; the PSF is substantially stable across the field. A standard analysis, analogous to the one described in the previous subsection, yields about 1000 stars; the reconstructed image is shown in Fig. 11, right. In this case the answer is not known, of course. The performance of the code may be evaluated by means of an experiment with artificial stars, represented by scaled copies of the PSF placed randomly in the field. We may generate




Figure 11. Left: image of the Galactic Center, observed with the PUEO system on the CFH Telescope. Right: reconstructed image, given by the superposition of ~ 1000 scaled copy of the PSF on a smooth background estimate. This reconstruction is actually the synthetic field produced by StarFinder.




Figure 12. Photometric and astrometric accuracy in the Galactic Center image, evaluated by means of an experiment with artificial stars (see text). Left: estimated vs. true magnitude for successfully detected artificial stars. The magnitude zero-point on the horizontal axis is arbitrary. Right: astrometric errors, represented by the distance between the true and the calculated position.




Figure 13. Left: image of the Trapezium cluster obtained with the ADONIS system on the ESO 3.6m telescope. Right: standard deviation of the magnitude measurements for the common stars in four exposures of the field. The comparison of the four lists has been performed with the algorithm outlined in Fig. 14. The magnitude zero-point on the horizontal axis is arbitrary. The error $\sim 0.3 \text{ mag}$ is associated to a star whose image, in two frames out of four, is corrupted by negative pixels.

for instance 10 frames, adding to the original image 10% of synthetic stars at random positions for each magnitude bin in the estimated luminosity function. The 10 frames must then be analyzed separately and the lists of successfully detected artificial stars have to be merged together. The final goal is to evaluate the photometric and astrometric errors for these stars. The results of a similar experiment are shown in Fig. 12.

4.3. Real data: internal accuracy with multiple frames

The last example is a set of four images of the Trapezium cluster (Fig. 13, left), taken in the K' band at slightly different pointings, though with a large overlap. The adopted observational strategy suggests us to compare the independent magnitude measurements for the corresponding stars in the four frames, in order to evaluate the internal accuracy of the reduction procedure. The results are shown in Fig. 13, right. The internal errors are generally below $\sim 1/10 \text{ mag}$, except one case for which the dispersion is higher. A careful analysis of this object indicates that its image is corrupted by a bunch of negative pixels in two frames out of four, perhaps due to some sky subtraction troubles in the calibration phase: this fact accounts for the larger error.

5. Possible improvements

5.1. PSF estimation

The PSF is basically estimated by combining the images of some suitable stars in the observed field: each pixel of the resulting PSF array is the median of the




Figure 14. Comparison of four lists of stars two at a time. Matching lists 1 and 2 produces two output lists 1' and 2' of common objects. The same for inputs 3 and 4 and corresponding outputs 3' and 4'. It may happen however that the objects in lists 1' and 2' form a different set from those in 3' and 4'. Therefore it is necessary to compare 2' and 3', for instance, to find a common set, represented by 2'' and 3''. Finally the intermediate lists 1' and 4' must be compared to 2'' and 3'' respectively. In this way the lists 1'', 2'', 3'' and 4'' should include the same set of objects.

corresponding pixels in the selected stellar images. The median is preferred to the mean because it is less sensitive to the so-called outliers (see Sect. 2.2.). Nevertheless it might be interesting to implement other strategies, for instance selecting not the median but the minimum value of the corresponding pixels. This operation may help estimating the PSF with few stellar images available possibly close one another, for examples the components of a crowded group.

5.2. Undersampled data

StarFinder has been designed for AO observations, which are commonly well-sampled. Tests on HST undersampled data have shown that its performance worsen when the Nyquist-sampling condition is not fulfilled; the results are slightly worse than those achievable by other methods, which use analytic models to represent the PSF. The obvious reason of this behaviour is that StarFinder matches the PSF template to a given star by numerical interpolation: the shifting errors due to the undersampling affect the astrometry and the photometry. In order to handle these cases it might be interesting to provide the program with a variety of PSF models that could be translated exactly with no interpolation. The PSF extraction phase would be replaced by a calibration of the most suitable model, accomplished by fitting some user-selected stars in the field. The rest of the code would remain substantially unchanged. Apart from the classical models given by superpositions of gaussians, Moffat functions, etc., it might be worthwhile to include templates which better reproduce the shape of a high-resolution PSF, for instance various Airy patterns degraded by static aberrations and partial AO correction.

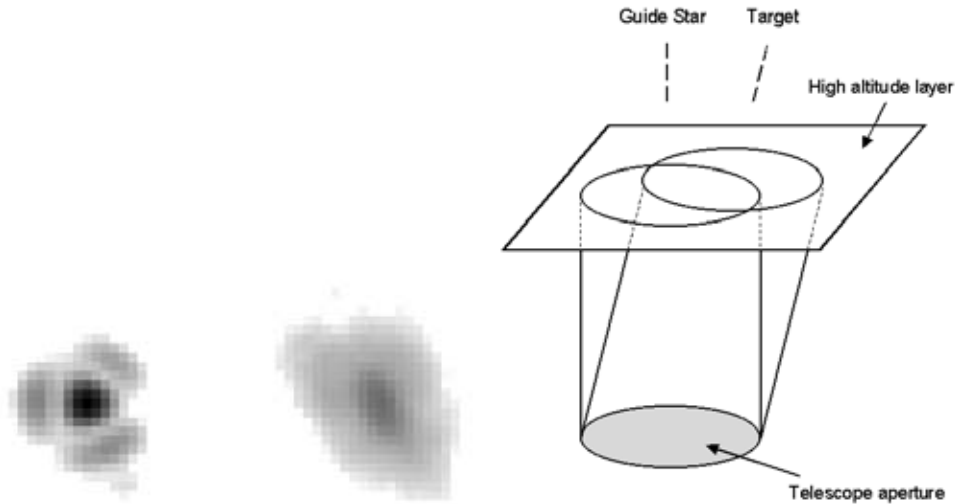


Figure 15. Left: image of the guide star in a $1' \times 1'$ observation taken with the AO system of the UH telescope (Simon, Close, & Beck 1999). Center: off-axis star (~ 40 arcsec away) in the same frame; the major axis of the stellar image points at the guide star. Notice the Strehl ratio reduction (a factor ~ 3) and the broadening (~ 2 times broader at the half-maximum level). Right: schematic explanation of the anisoplanatic effect. A high altitude layer is represented, along with the footprints of the wavefronts coming from the guide star and the off-axis source, as seen from the telescope aperture. The footprints do not perfectly overlap; therefore the correction, optimized for the guide star, is not optimal for the off-axis source. The linear overlap of the footprints decreases more quickly in the radial direction as the angle between the two beams increases: this is the reason of the radial elongation.

5.3. Anisoplanatism

The last possible improvement discussed here is related to the anisoplanatic effect, i.e. the PSF variation across the image (see Fig. 15). In the case of classical AO (one star as reference), the effect can be described in terms of a blurring and radial elongation of the off-axis PSF, with respect to the on-axis guide star.

In general the analysis of an anisoplanatic field requires the knowledge of the local PSF. A first strategy is to partition the frame and estimate a PSF for each sub-region by combining some suitable stellar images. There are some difficulties related to this approach. In crowded fields it might be very difficult to extract a reasonable PSF estimate for each sub-region, due to the overlapping of the stellar images. In uncrowded fields there may be sub-regions with too few stars to estimate the PSF. Moreover there might be a differential zero-point effect in the magnitude scale between different sub-regions, arising from the uncertainties in the PSF normalization already discussed in Sect. 3.5..

Another solution is based on the possibility to represent the off-axis PSF as the convolution of the guide star with an anisoplanatic function (Voitsekhovich et al. 1999, Fusco et al. 2000). A method to reconstruct the anisoplanatic kernel has been described by (Fusco et al. 2000): this technique relies on the knowledge of the correction introduced by the AO system and the statistical properties of the atmospheric turbulence. An alternative solution is to represent the anisoplanatic function with some parametric model, accounting for the PSF blurring and elongation; the spatial dependence of the parameters may be calibrated by fitting some stars in the field.

A significant reduction of the anisoplanatism will be possible in the near future using Multi-Conjugate Adaptive Optics (Beckers 1989, Ellerbroek & Rigaut 2000, Ragazzoni, Farinato & Marchetti 2000), a technique still under development which uses more than one reference star in order to have an acceptable correction over a wider field of view. Even in this case, however, some anisoplanicity is to be expected, perhaps with more complicated patterns.

Acknowledgments. Emiliano Diolaiti is grateful to Andreas Quirrenbach for the invitation and to Vesa Junkkarinen for distributing the code to the participants. Orazio Bendinelli, Gianni Bernardi, Domenico Bonaccini, Laird Close and Gianluigi Parmeggiani are acknowledged for their contribution to the development and testing of StarFinder. Thanks are due to François Rigaut for kindly providing the image of the Galactic Center and for supporting the initial development of StarFinder, and to Roberto Ragazzoni and Massimiliano Tordi for useful discussions on AO and anisoplanatism.

References

- Almozino, E., Loinger, F., Brosch, N. 1993, MNRAS, 265, 641
 Beck, J. V. , Arnold, K. J. 1977, Parameter estimation in engineering and science, John Wiley & Sons
 Beckers, J.M. 1989, in SPIE Proc. 1114, 215
 Bendinelli, O., Parmeggiani, G., Piccioni, A., Zavatti, F. 1987, AJ, 94, 1095
 Bertin, E., Arnouts, S. 1996, A&AS, 117, 393
 Bevington, P. R., Robinson, D. K. 1994, Data reduction and error analysis for the physical sciences, McGraw-Hill
 Bijaoui, A. 1980, A&A, 84, 81
 Diolaiti, E., Bendinelli, O., Bonaccini, D., Close, L. M., Currie, D. G., Parmeggiani, G. 2000, A&AS, accepted (astro-ph/0009177)
 Ellerbroek, B.L., Rigaut, F.J. 2000, in SPIE Proc. 4007, 1088
 Esslinger, O., Edmunds, M. G. 1998, A&AS, 129, 617
 Fusco, T., Conan, J -M., Mugnier, L. M., Michau, V., Rousset, G. 2000, A&AS, 142, 149
 Gonzalez, R. C., Woods, R.E. 1992, Digital image processing, Addison-Wesley
 Landsman, W.B. 1995, in ADASS IV Proc., ed. R.A. Shaw, H.E. Payne, & J.J.E Hayes (San Francisco: ASP), 77, 437

- Luenberger, D. G. 1984, *Linear and nonlinear programming*, Addison-Wesley Publishing Company
- Park, S., Schowengerdt, R. 1983, *Computer Vision, Graphics & Image Processing*, 23, 256
- Ragazzoni, R., Farinato, J., Marchetti, E. 2000, in *SPIE Proc.* 4007, 1076
- Rigaut, F., Salmon, D., Arsenault, R., Thomas, J., Lai, O., Rouan, D., Véran, J. -P., Gigan, P., et al. 1998, *PASP*, 110, 152
- Simon, M., Close, L.M., Beck, T.L. 1999, *AJ*, 117, 1375
- Sosin, C., King, I.R. 1997, *AJ*, 113, 1328
- Stetson, P. B. 1987, *PASP*, 99, 191
- Véran, J. -P., Rigaut, F., Maitre, H., Rouan, D. 1997, *J. Opt. Soc. Am. A*, 14(11), 3057
- Véran, J. -P., Rigaut, F. 1998, in *SPIE Proc.* 3353, 426
- Voitikhovich, V.V., Bara, S. 1999, *A&AS*, 137, 385