Pattern and Transient Removal Algorithms[§]

Francisco G. Valdes

NOAO[*], P.O. Box 26732, Tucson, AZ 85726, USA

ABSTRACT

This paper describes algorithms for the removal of two types of contaminating features in astronomical images. One is background patterns that are spatially fixed (over many exposures) but variable in amplitude. These include fringing and pupil patterns. An algorithm for automatically fitting the pattern amplitude is presented. The second type of contamination consists of non-astronomical (cosmic rays and satellite trails) and astronomical (asteroids) transient sources. An algorithm for detecting and removing these sources from stacked images is also described. Implementations of these algorithms in IRAF are illustrated with data from the NOAO large format mosaic imagers taken by the NOAO Deep-Wide Field Survey (NDWFS)[†].

Keywords: image processing, pattern removal, transient detection, IRAF, ACE

1. INTRODUCTION

This paper discusses two related topics concerning the removal of contaminating features in astronomical images. The first is the removal of background patterns where the amplitude of the pattern is scene dependent. Patterns of this type include fringing and reflections producing an out of focus pupil pattern. These patterns, defined by calibration template images, are removed by scaling the pattern amplitude to match a data image and then subtracting from or dividing into the data image. The challenge is to automatically determine the amplitude that best removes the pattern from each observation. This paper describes an algorithm that robustly determines the scaling.

The second type of contaminating features are transient sources affecting stacks of overlapping exposures. The transient sources include cosmic rays, satellite trails, and asteroids. Some of these sources have interesting astronomical significance but in the context of making deep static images of the sky from many exposures these are contaminants to be removed. In this paper we describe some early work on this problem.

A key aspect of both the pattern and transient removal algorithms is the treatment of the astronomical scene which, in this context, is a source of interference. For pattern removal this consists of avoiding pixels containing the astronomical objects using a *masking* technique. One must also *map* the non-pattern background in order to isolate the pattern to be removed. Transient removal involves *image subtraction* to cancel out the non-transient scene. These are large and significant algorithmic areas which cannot be covered in this paper. Instead, we briefly mention the tool used to provide these features in the implementation of the removal algorithms presented below.

The IRAF *Astronomical Cataloging Environment* (ACE)[2,3] is a software package being developed for detecting, classifying, and cataloging astronomical objects in images. One way ACE differs from other packages is the

E-mail: fvaldes@noao.edu

*National Optical Astronomy Observatories, operated by the Association of Universities for Research in Astronomy, Inc. (AURA) under cooperative agreement with the National Science Foundation.

†This work made use of images provided by the NOAO Deep Wide-Field Survey (Jannuzi and Dey[1]) supported by the National Optical Astronomy Observatory (NOAO).

‡Image Reduction and Analysis Facility, distributed by the National Optical Astronomy Observatories.

inclusion of special features beyond object detection and cataloging which may be used in algorithms such as those described in this paper. The features which are applicable to the algorithms are the production of *object masks* and *background maps*, and built-in *image difference detection*. The algorithms and concepts are described in Ref. 2.

Due to the press of other projects, the `ACE` package has not been publicly released. However, the object and background identification capabilities have been released as a limited, though useful, application that produces masks of the objects and maps of the background. Both of these are in space efficient formats understood by other IRAF applications. In particular, objects are identified in a compact run-length encoded mask format used by many IRAF applications and the background is output as a lower resolution image that applications can expand as needed. This task is called `OBJMASKS` and appears in the latest version of IRAF.

In algorithm papers, such as this, there is a mixture of algorithmic ideas and implementation ideas. The implementation is always in some specific environment. Here the example implementations are in IRAF[‡] and also make use of some general IRAF concepts such as *pixel masks*. But the reader should keep in mind that this is not a paper about specific programs but about the generic algorithmic concepts behind them. These may be implemented in other environments.

2. REMOVING SCENE DEPENDENT INSTRUMENTAL PATTERNS

There are many instrumental patterns that appear in astronomical exposures. The most common are additive or multiplicative patterns. These are removed by subtracting or dividing a template of the pattern from the observations. The pattern template is generally derived from calibration exposures or from the set of data exposures. When the pattern is independent of the field of view and independent or proportional to the exposure time the pattern template can be simply subtracted or divided from the astronomical observations without modification or with just a numeric scaling based on the exposure times. The typical scene independent patterns of this type are bias, dark count, and flat fields.

However, some patterns vary in amplitude from exposure to exposure. The variation is caused by differences in the field of view. In other words the pattern amplitude is scene dependent. Examples of this are CCD fringing and defocused pupil patterns. In the first case the fringing varies with the strength of the atmospheric emission lines and in the second on the amount of light in, or even outside, the field of view.

These patterns require determining a scale by which the amplitude of the template pattern is adjusted to match the amplitude in each observation. This is essentially a simple least squares problem except that the astronomical scene and noise "contaminates" the observation. This section presents an algorithm for automatically determining the scaling for every observation in the presence of these contaminants. This algorithm must be robust in order to be used with confidence on large datasets in automated pipelines or batch processes.

Figure 1 shows an observation containing both a pupil and fringe pattern. This example was selected as a particularly challenging case where bright sources contaminate the pupil pattern in one quadrant.

In the next section we describe an algorithm for determining the scale factor between a pattern template and a data image. How the the pattern template is determined is not discussed in this paper. This may be found in the description of mosaic CCD reductions, which is equally valid for single detectors, given by Valdes.[4] That paper also gives more detail on pupil and fringe removal. Note, however, that it was written before the development of the automated pattern fitting algorithm presented here.

2.1. Pattern Fitting Algorithm

The object of the pattern fitting algorithm is to determine the scaling factor that matches a pattern template image to a data image in the presence of non-pattern sources and noise. If the pattern in the data is stronger than the astronomical sources or the noise then it is easy to remove. However, the common situation is where the pattern is very weak. The signal in individual pixels may be much smaller than the pixel noise but is evident because of its large scale coherence. So the problem can be described as finding and removing a weak coherent pattern in the noise of the background.

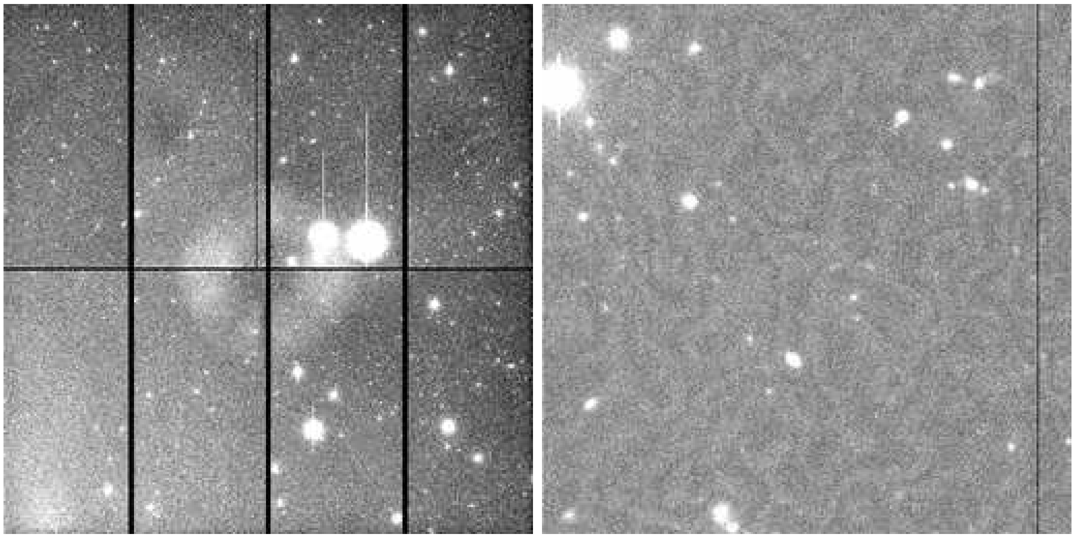

Figure 1. Example of pupil and fringing patterns in an observation with bright and faint objects, a non-flat background, and noise. The data is from the NDWFS using the NOAO Mosaic Imager at the Kitt Peak National Observatory's 4-meter Mayall Telescope. The left panel shows the entire mosaic field of view (8192x8192 pixels, 36x36 arcmin, 0.25 arcsec/pixel). The large doughnut in the middle is an out of focus pupil pattern caused by a small fraction of the light reflected between the filter and corrector. The panel on the right is a small 512x512 portion of the same observation showing fringing. The fringing extends over the entire set of detectors.

There are two key concepts to the pattern fitting algorithm. Both are natural to the way a human observer decides on the best fit. The first, and most important, is to ignore the astronomical sources in the field and concentrate on the regions containing the pattern. Note that the astronomical sources means not only objects, such as stars and galaxies, but also the background sky. The pattern is generally a modulation on some background and so the background must also be identified and ignored, which means considering the amplitude of the pattern relative to the background. Ignoring such features easy and automatic for humans but a challenge for software.

A simple first approach that was tried to eliminate or minimize the effects of objects was to use sigma clipping. This did not work very well for several reasons. One important factor is that extended or bright objects have faint extended light that always falls below any clipping threshold and significantly affects the scale determination in weak patterns.

The second aspect of what an astronomer does in deciding on the best fit is to preferentially look at where the pattern is most obvious. In algorithmic terms this means that greater weight needs to be given where the pattern deviates from the background.

These two considerations lead to the following algorithmic features. The astronomical objects must be detected, including as much of the extended faint light as possible, and "masked" from the scale determination. The algorithm must also be told where the pattern is localized if it does not extend across the full field of view. These are related by considering both as masking regions that do not contain useful information about the pattern amplitude. The background in both the pattern and data must be determined and accounted for when the pattern is a variation on top of the background. Finally, the calculation of the scale using the remaining data should apply weights defining where the pattern is most important. While the obvious choice of weights is the pattern itself, we generalize this by allowing users to define the weights.

Based on these considerations, the pattern fitting algorithm determines the scaling that minimizes the weighted mean difference between an input image and a pattern image as expressed in Eq. 1.

$$\langle((A-B)-\alpha(P-Q))(W-V)\rangle_M = 0 \tag{1}$$

where α is the fitting parameter to be determined. The other terms are the data image containing the pattern to be fit A, the background in the data image B, the pattern image to fit to the data image P, the background in the pattern image Q, a weight image W, and the background in the weight image V. The mean is computed over a subset M of the image pixels.

In this general formulation we have allowed all the quantities to be functions of position, i.e. images, and to be different. In practice, some of these images may be absent, constant, or the same as one of the other images. One common situation is to set the background images to be the mean of the image to which it applies. In this case the statistic becomes

$$\langle((A-\langle A\rangle)-\alpha(P-\langle P\rangle))(W-\langle W\rangle)\rangle_M = 0 \tag{2}$$

The places where the pattern image deviates most from the pattern background is what is most noticeable to the eye. Therefore the typical case for weighting is to use the pattern itself. For this case the statistic is given by

$$\langle((A-B)-\alpha(P-Q))(P-Q)\rangle_M = 0 \tag{3}$$

A final variant, which has been found effective, is to use a smoothed version of the pattern image for the weighting. This is advantageous when the pattern image has low signal-to-noise and the pattern does not include high spatial frequencies. A moving average has proven to be a good and efficient filter that can be evaluated at the same time as the statistic is evaluated.

The mean value in these equations is computed over all the pixels in the data image A containing the pattern. As noted earlier, a key to obtaining a good scale is to restrict the calculation to background pixels which contain the pattern. Algorithmically we identify the pixels by the region denoted by M. One may think of the region as a mask where only those pixels which pass through the mask are used. The mask may actually be composed of the overlap of several masks. Logically one has a mask blocking the astronomical objects and another mask blocking the regions where the pattern contains no signal. A reason why these might be separated is that the object mask must be determined for each data image while the pattern mask remains the same for all data using the same pattern template.

One may solve eqs. 1, 2, or 3 in various ways. Here we point out one way that illustrates how to do the computation efficiently and flexibly. Consider Eq. 1. Associate all the multiplicative terms and use the linear property of the mean to obtain

$$\alpha = \frac{\langle AW\rangle - \langle AV\rangle - \langle BW\rangle + \langle BV\rangle}{\langle PW\rangle - \langle PV\rangle - \langle QW\rangle + \langle QV\rangle} \tag{4}$$

where we have dropped the subscript M which is understood to apply to each mean. Note that in this fraction the number of pixels is the same in both the numerator and denominator so the means may be replaced by the sums of each quantity over the pixels in M, though the normalization by the number pixels in M is not a major operation.

When some of the quantities are absent, as for example when no backgrounds are present, then some of the terms may be eliminated. When some of the parameters are constants the sums of the remaining quantity may be computed and the constant multiplied at the end to save a multiplication at every pixel. Also consider what happens when two images, namely the pattern and weight images, are the same. For instance, let $W = P$, $V = Q$, and B and Q be constant. Then the statistic reduces to

$$\alpha = \frac{\langle AP\rangle - Q\langle A\rangle - B\langle P\rangle + BQ}{\langle P^2\rangle - QQ} \tag{5}$$

Another interesting result is that when a mean background is selected, $B = \langle A \rangle$ and $Q = \langle P \rangle$, the problem reduces to

$$\alpha = \frac{\langle AP \rangle - \langle A \rangle \langle P \rangle}{\langle P^2 \rangle - \langle P \rangle^2} \tag{6}$$

for the case in Eq. 5. As noted earlier, using a mean for the backgrounds involves nothing more than accumulating certain sums in a same single pass through the image.

Thus far we have not discussed the meaning of the backgrounds allowed in Eq. 1 and its relatives. The derived scale depends strongly on matching the data and pattern backgrounds. In particular, the data image background and the pattern background must correspond to the same feature of the pattern. So if the mean of the pattern is used as the background then the data background must correspond to the mean of the pattern in the data. If the pattern has a zero background outside the pattern and zero is specified for the pattern background then the data background must be that unaffected by the pattern in the data. In general, if one uses masks to isolate the regions of the pattern and to exclude scene objects then the mean backgrounds are appropriate. The exception to this is if the pattern is not localized and there is a background gradient in the input data which is not part of the pattern. In that case an input background image, such as provided by the background map produced by `ACE`, can be used while the mean of the pattern template is used for the pattern background.

2.1.1. Mosaic camera observations

Mosaic cameras tile the focal plane with detectors and produce a set of images, one or more from each detector, for each exposure. Focal plane patterns, such as the pupil, may extend across all the images. Detector patterns, such as fringing, may also occur in all the detectors.

There are several ways the patterns can be removed. One is simply to treat each image independently. However, generally pupil patterns and fringing should scale by the same amount across the field of view since they vary by the amount of sky or field light illuminating the telescope. To determine a single scale factor one could pack the exposures into a single images. This should be done without interpolation degrading the resolution. If the data acquisition system produces this type of format it would be sensible to apply the pattern fitting algorithm to this image, taking into account any gaps in the masks.

However, it is more common to keep the individual exposures separate. This has various advantages. A nice feature of the pattern fitting algorithm is it can be applied easily to this case. All that is required is that the accumulation of the various statistics be continued across all the pieces. Naturally, each image should use its own mask to select significant data. At the end of the accumulation stage the scale statistic α, which applies across the whole mosaic exposure, can be calculated as with a single image.

Note, it is also easy to maintain separate sums for each piece in addition to the global sums and thereby obtain scale factors determined independently for each piece. It is then possible to use the average of the independently derived scales as another estimate of the global scale factor or use the variation of the scale factors as a diagnostic.

One may choose to include only a subset of the mosaic images in computing the global scale. This would be either because the pattern does not impact all the pieces or a subset of the data, even a single image, may be sufficient to accurately determine the scale.

2.2. Identifying Objects and Background

Two key aspects for accurately determining the scale factor in this algorithm is selecting the pixels defining the region M and accounting for background gradients. These two topics are opposite sides of the same coin. To segment an image into objects above a background requires determining the background uninfluenced by the objects.

In our implementation we use the IRAF task `OBJMASKS` to produce masks and background maps. The algorithmic details of this are given in Ref. 2. Figure 2 shows the application of this task to the data from

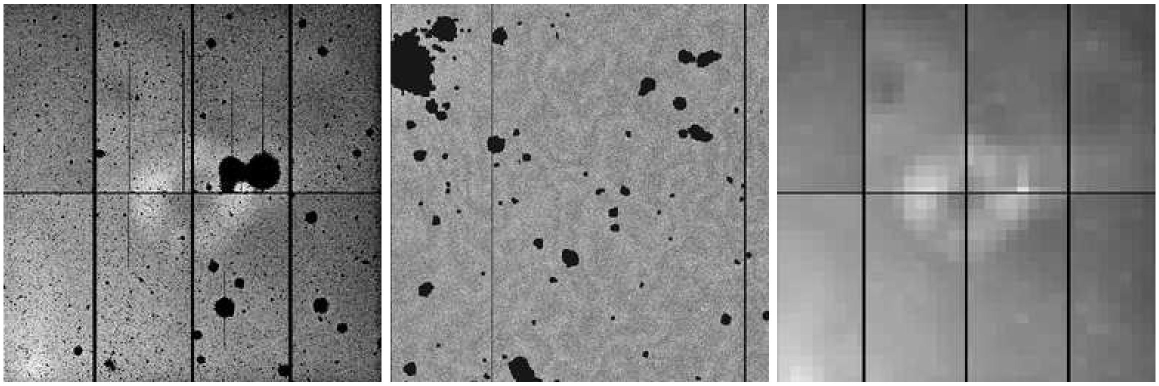

Figure 2. Example of object masking and background determination. The right two panels are for same data as Fig. 1. In both the object mask produced by the task `OBJMASKS` is overlayed. In this black and white reproduction the mask pixels are shown as black but typically the overlay can be color coded to discriminate bad pixels, saturated and bleed pixels, and the objects. The background image produced is shown in the third panel. The lower resolution sky shows, in this case, the pupil pattern which has not yet been removed and a brightening of the background towards the lower left of the mosaic field. The other faint ring-line structures are residual flat field patterns that are removed by a sky flat produced after the pupil and fringe patterns are removed.

Fig. 1. The images show overlayed object masks; the object fluxes are not modified in the images but are simply blocked by the overlayed mask in the display. Also shown is the background map produced by the task.

One feature of the object mask creation which permits handling of the extended light around objects is the *growing* of the object regions. The growing may be specified as a combination of some number of rings around the irregular shape of the detection and adding rings until the total area of the object region is some factor, such as twice, that of the original area.

In order to fully automate the pattern removal algorithm presented here the object masks must also be automated. `ACE` can be used to automatically produce object masks for each data image either in a pipeline script or as a batch process run over a night's set of observations. Only in cases with confused or little background would adjustment of the default parameters be needed.

The usefulness of creating objects masks extends beyond pattern removal. There are other algorithms which can benefit from such masks. Examples of these are determination of sky flat field and fringe pattern templates using unregistered exposures. Sigma clipping is sometimes used but this typically leaves residuals from the extended halo light of objects which are below the sigma threshold. Examples of creating these templates using image combining with object masks are given in Ref. 5.

2.3. PATFIT

The implementation of the pattern fitting algorithm in IRAF is a task called `PATFIT`. This task is very general and includes options for a variety of types of removal. It also provides for all the combinations of backgrounds, weights, and mask images and for mosaic data. For convenience this general task is usually called from simple scripts specialized for certain applications. In particular, there are separate scripts specialized for pupil (`RMPUPIL`) and fringe removal (`RMFRINGE`) in multiextension mosaic data. In this section we describe the various features of this implementation.

First we consider what happens after the scaling is determined. While one result of the task is simply to report the scaling factor and other documentary information, its goal is generally to remove the pattern. Since the task already has the images it makes sense to have it also produce corrected output images.

Table 1 shows the output types from `PATFIT`. These define the fit including all the background terms. In other words, the fit is what transforms the pattern template to match the data image. The output may then be the fit (*fit*), the difference or residual (*diff*), or the ratio (*ratio*). The *flat* output type normalizes the fit by the

Table 1. Definitions of the `PATFIT` Output Types

fit:	$\alpha * (P - Q) + B$	
diff:	$A - (\alpha * (P - Q) + B)$	
ratio:	$A/(\alpha * (P - Q) + B)$	
flat:	$A * B'/(\alpha * (P - Q) + B)$	where $B' =< A - \alpha(P - Q) >$
	$= A/(\alpha' * (P - Q) + B/B')$	where $\alpha' = \alpha/B'$

mean of the data with the fit subtracted. This last type is used to correct flat fields while maintaining values comparable to the original flat field counts rather than values near unity.

In many applications what is actually desired to to remove only the part of the fit related to the amplitude variable component of the pattern. There are variations of the four output types which set B and/or Q to zero during output though the terms are still used in the fit. For both pupil pattern and fringe removal normally B would be set to zero during output to leave the background in the output. Whether Q is included or not in the output depends on how the pattern template is created and fit. For pupil pattern templates that are created with a zero background outside the pattern, but where the fitting uses $B = \langle A \rangle$ and $Q = \langle P \rangle$ within the region of the pattern, Q would also be set to zero in the output.

The task input (data, pattern, backgrounds, and masks) may be single images or mosaic multiextension FITS files. The mosaic input may be restricted to a subset of pieces for determining the pattern scale factor but the pattern removal is performed over all the pieces. The various terms may be lists to allow operation over a large data set. Because the algorithm has proven to be very reliable it is common to operate on a whole night's data as a batch process using a list of input images, input masks, and output files but with a single pattern template for the night.

The backgrounds may be full images, maps (lower resolution images such as produced by `ACE` and shown in Fig 2), constant values, or a special option to use the mean of the image to which it applies. The task has been carefully designed to recognize and efficiently handle all combinations of constants and when the weight and pattern image are the same. The case of a zero constant value is also specially handled.

An option in the task is to apply a boxcar smoothing to the weight image. For the case of the pupil pattern in the NOAO Mosaic data the smoothing is required to get the best result.

The exclusion of pixels containing objects and in regions where the pattern is absent is a key component of the algorithm. The pixels to exclude are specified by IRAF *pixel masks.* It was noted earlier that this can be done with one mask but may also be done with two masks, one for the objects in the data images and one for the pattern images.

In addition to the corrected data, the task produces diagnostic information and processing keywords in the output data. The processing keywords allow determining if a pattern has been removed, what pattern template was used, and what scale was determined and removed.

Figures 3 and 4 show the results of automatically applying `PATFIT` to the example data shown in earlier figures. Fig. 3 illustrates removal of the pupil pattern and Fig. 4 illustrates fringe removal. Both figures also overlay the object masks. Using the objects masks when displaying the data can be useful in being able to study the background without distractions from the objects. Also, in the IRAF display task the object masks can also be used to automatically compute the best grayscale display range for the background.

A comparison has been made between the result of an astronomer interactively adjusting the scale factors and the automatic result of the algorithm. As mentioned earlier, the production of the objects masks for each data image using `ACE` is also an automatic process and the reliablity of this was similarly checked. The comparisons and checks were done over a few nights of data. In all cases the astronomer felt the algorithm did as well or better than they could interactively.

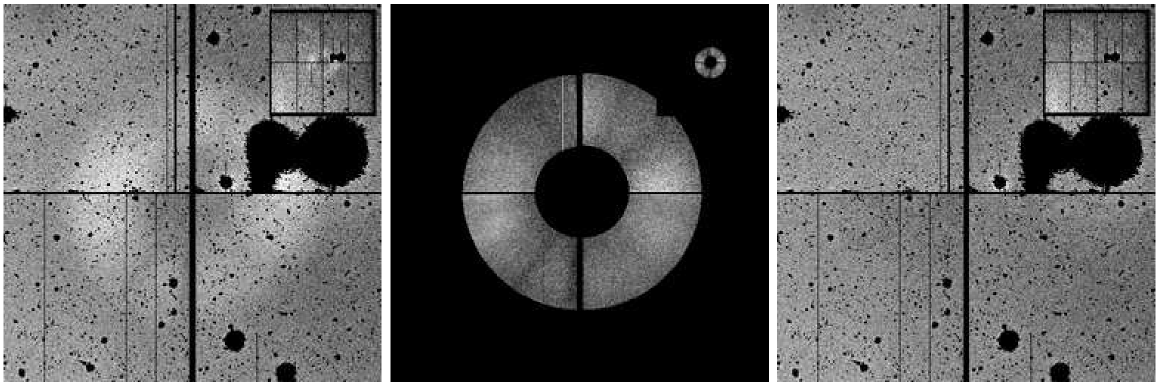

Figure 3. Example of pupil pattern removal from the same NDWFS data shown earlier. The three panels are to the same scale. The left panel shows the pupil pattern visible in the data exposure with the objects masked (the black regions excluding the central cross corresponding to the gaps between the CCDs). The pattern covers only parts of the central four CCDs in the mosaic. Only these CCDs are used in the computation of the pupil scaling. The central panel shows the pupil pattern template derived from the night's data. The pattern template has been overlayed by a circular mask defining just the area of interest for the pattern as discussed in the text. The right panel shows the same data as in the left panel but with the pupil pattern automatically scaled and subtracted. All these figures include a view of the whole mosaic in the overlayed subwindow at upper-right corner which should not be confused with the actual data.

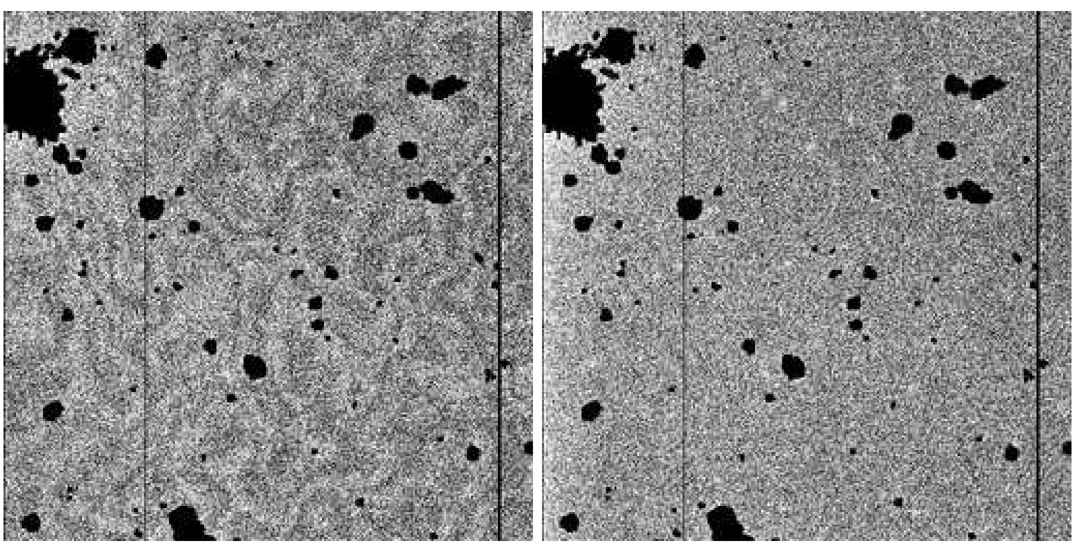

Figure 4. Comparison of before and after fringe removal. The black regions are masked objects and cosmetic defects. The fringe scaling from a template, created by stacking many unregistered exposures during the night, was automatically determined and subtracted. The subtraction is optimized over the entire mosaic field.

3. REMOVAL OF TRANSIENT SOURCES

Transient sources are, by definition, things which appear in only one or a few exposures or vary in position, brightness, and/or shape between exposures. These may be categorized as non-astronomical, such as cosmic rays, satellites, and airplanes, and astronomical, such as asteroids, variables, and supernovae. The astronomical sources are often the target or interesting by-product of observations which one wants to extract.

However, in programs where deep images of the static sky are created by *stacking* multiple, possibly dithered, exposures, the transient sources are contaminants. These must be identified and removed from each exposure before stacking. The two most vexing sources are cosmic rays and satellite trails. Note that in wide-field imaging with large format detectors the occurrence of satellite trails is fairly common.

The algorithm for removing the transient sources in these types of programs is to create a first pass deep stacked image using a fairly heavy pixel sigma clipping to largely eliminate the transients. This is not a good algorithm for the final deep image because sigma clipping adversely affects the PSFs of brighter objects and it fails to remove the fainter wings of sources such as satellite trails that fall below the statistical threshold in a pixel but are still visible as coherent patterns.

However, this image is a good reference for subtracting from each individual exposure. This differencing step plays the important role of eliminating the interfering effics of the static astronomical source. Note that it is possible to substitute some other image provided it has comparable or better depth and resolution.

Detection of features in the difference identifies the transient sources. Faint source detection algorithms are specifically optimized for just the problem of finding coherent regions where individual pixels may be within the statistical noise but objects can be identified as extended coherent regions. In addition, once the regions are found they can be extended by growing the boundaries to eliminate even fainter light associated with the objects.

The regions containing the transient sources are then excluded when stacking the images to form a final deep image. The final stack is created without pixel clipping to avoid the effects on the PSF noted earlier.

Implementation of this algorithm can be broken down into three steps, each with their own complex subalgorithms. One is registration and subtraction of two images. Registration, in this and the stacking step, involves matching positions, fluxes, and PSFs. The second step is the detection of transient sources in the difference. This must also include elimination of artifacts from errors in the subtraction. The last step is registration and stacking of images. This step is used twice, once with sigma clipping to form the reference image and once using information about the sources from the detection step.

This algorithm was implemented for creation of the final deep images in the NDWFS. The first two steps are performed using the `ACE`[2] package. The two steps are combined into a single step. In other words, rather than producing a difference image as a separate step, the program takes the data and reference images as input and does the difference as the detection takes place. This includes matching the images spatially and in flux. It also includes computing the background and background noise in the separate images and then using these to define background and noise in the difference. What is currently not included is matching of the PSFs. This is planned using algorithms pioneered by Alard.[6]

In the basic detection stage the obvious noise detections are eliminated as early as possible. A common example is requiring a minimum number of pixels to be associated with a detection. For the special case of detection in difference images, and a major reason why it is useful to work from the two separate images and have the program do the difference on the fly, is that then there can be some basic tests done involving the original data. In `ACE` the flux in the source detected in the difference image is compared to the flux in the reference image. In simple terms, detections in the difference which are associated with strong underlying sources are identified and excluded. Depending on the ability of the difference step to match PSFs, the most common artifact in the difference image are residual cores of bright stars due to mismatched point spread functions.

As note earlier, the source detections performed in `ACE` produce object masks including growing to identify additional low level light. The object masks provide the connection to the final stacking step.

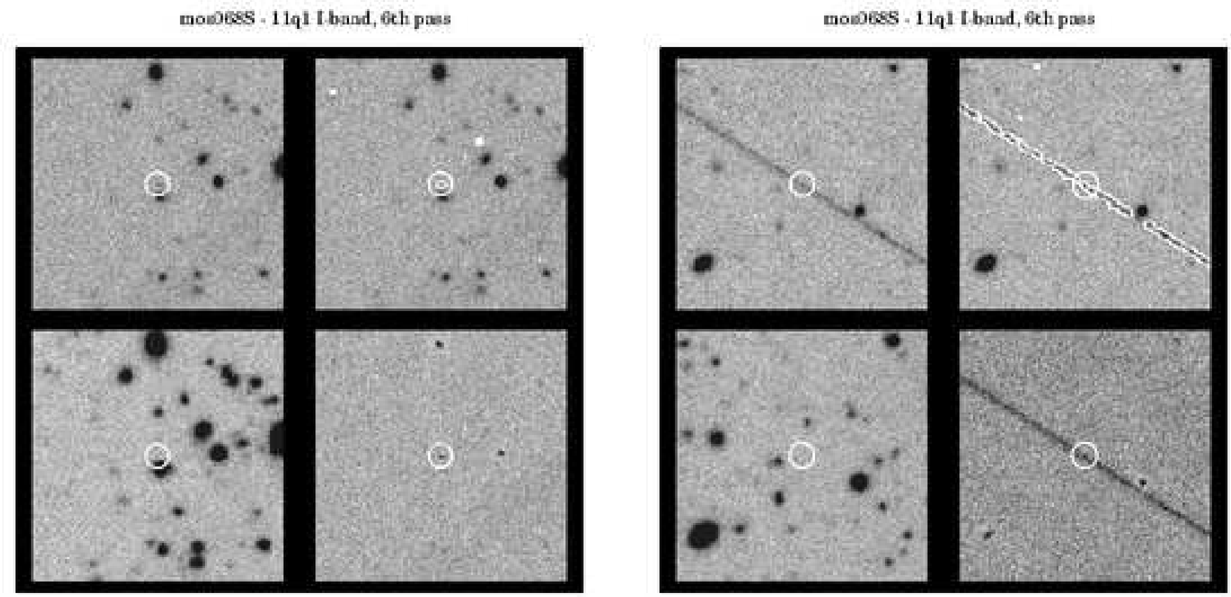

Figure 5. Examples of transient object detections using differencing between a single exposure and a clipped stack. The panel on the left shows detection of a cosmic ray and the one on the right of a satellite trail. Each panel has four panes which show matching cutouts centered on a detection. The upper-left is the single exposure with the candidate source marked. The upper-right is the same but with the detection "isophote" included. The lower-left is the matching view in the reference image with the location of the detected source overlayed for comparison. The lower-right is the difference image. One thing to point out is that in the difference image there are some other sources which are generally residuals of brighter sources in the individual images. These are generally excluded from detection by a flux comparison criterion.

The image stacking step, both to form the reference image with sigma clipping and the final image with object masks, is implemented using the IRAF task `IMCOMBINE`.

This algorithm for detecting and removing transient sources can be performed non-interactively. However, in the current implementation there are artifacts, such as scattered light, that are not well handled. So the astronomer normally reviews the detected sources. Figure 5 shows displays from this review which illustrate the method.

The current implementation is still fairly primitive. There are a number of enhancements, in addition to classifying and cataloging the sources as serendipitous science, that can be made to improve the removal of the contaminating sources. One is to use better PSF matching and image subtraction algorithms,[6] though often the seeing is similar enough that the simple subtraction works quite well for this application provided the cores of stars are identified as described earlier. The reference image should be matched to the individual unresampled exposure to eliminate resampling effects in the image with the potential cosmic rays and transients. For the NDWFS mosaic data the image difference is currently done after all the images have been resampled to a common pixel grid. Another improvement is to do some intelligent adjustment of the difference detection object masks. In particular, recognizing the long streaks of satellite trails or fast moving asteroids which are sometimes broken up by an threshold detector and connecting the mask together and widening the mask perpendicular to the trail for best suppression of the profile wings in the final stacking.

REFERENCES

1. B. T. Jannuzi and A. Dey, "The NOAO Deep-Wide Field Survey," in *Photometric Redshifts and the Detection*

of High Redshift Galaxies, R. Weymann, L. Storrie-Lombardi, M. Sawicki, and R. Brunner, eds., *ASP Conf. Series* **191**, pp. 111–116, 1999. (http:www.noao.edu/noao/noaodeep/).

2. F. G. Valdes, "Ace: Astronomical Cataloging Environment," in *Automated Data Analysis in Astronomy*, H. P. Singh, R. A. Gupta, and C. A. L. Bailer-Jones, eds., pp. 15–29, Narosa, New Dehli, 2001.
3. F. G. Valdes, "Ace: Astronomical Cataloging Environment," in *Astronomical Data Analysis Software and Systems X*, F. R. Harnden, Jr., F. A. Primini, and H. E. Payne, eds., *ASP Conf. Series* **238**, pp. 507–510, ASP, San Francisco, 2001.
4. F. G. Valdes, "The reduction of CCD mosaic data," in *Automated Data Analysis in Astronomy*, H. P. Singh, R. A. Gupta, and C. A. L. Bailer-Jones, eds., pp. 309–353, Narosa, New Dehli, 2001.
5. F. Valdes, "New mosaic reduction capabilities." http://iraf.noao.edu/projects/mosaic/objmasks/.
6. C. Alard, "Image subtraction using a space-varying kernel," *A&A Supp.* **144**, pp. 363–370, 2000.