

Faint source detection in ISOCAM images

J.L. Starck, H. Aussel, D. Elbaz, D. Fadda, and C. Cesarsky

DAPNIA/SEI-SAP, CEA-Saclay, F-91191 Gif-sur-Yvette Cedex, France

Received December 2, 1998; accepted June 4, 1999

Abstract. We present a tool adapted to the detection of faint mid-infrared sources within ISOCAM mosaics. This tool is based on a wavelet analysis which allows us to discriminate sources from cosmic ray impacts at the very limit of the instrument, four orders of magnitudes below IRAS. It is called PRETI for Pattern REcognition Technique for ISOCAM data, because glitches with transient behaviors are isolated in the wavelet space, i.e. frequency space, where they present peculiar signatures in the form of patterns automatically identified and then reconstructed. We have tested PRETI with Monte-Carlo simulations of fake ISOCAM data. These simulations allowed us to define the fraction of remaining false sources due to cosmic rays, the sensitivity and completeness limits as well as the photometric accuracy as a function of the observation parameters.

Although the main scientific applications of this technique have appeared or will appear in separated papers, we present here an application to the ISOCAM-Hubble Deep Field image. This work completes and confirms the results already published (Aussel et al. 1999).

Key words: methods: data analysis — infrared: galaxies — cosmology: observations — methods: image processing

1. Introduction

Following the detection of ultra-luminous infrared galaxies (ULIRG's) by the IRAS satellite (Houck et al. 1984; Houck et al. 1985; Soifer et al. 1984a; Soifer et al. 1984b), it is not clear whether such objects, which are very bright but not numerous in the nearby universe, could be representative of a more common phase in the evolution of normal galaxies. In other words, could we expect the lack of detection of primeval galaxies to be due to dust extinction in systems emitting more than 90% of their light in

the infrared, as in local ULIRG's (Djorgovski & Thompson 1992).

Several programs were devoted to this search using ISOCAM (Cesarsky et al. 1996), one of the four instruments on board of the ISO (Infrared Space Observatory) spacecraft (Kessler et al. 1996) which ended its life in May, 1998. The present paper is devoted to reduction of data obtained with the long wavelength (LW) detector of ISOCAM, a (32×32) pixel array of SiGa. The LW detector operates in the range 4 to 18 μm , with a sensitivity four orders of magnitude better than IRAS and a spatial resolution sixty times better. This channel of ISOCAM was suited for the search of MIR dust emission in galaxies of redshifts typically below $z = 1.5$ (Elbaz et al. 1998). In this wavelength range, we find UIBs (Unidentified Infrared Bands) from 6.2 to 12.7 μm and Very Small Grains above 10 μm (Vigroux et al. 1998). However, because the 32×32 pixels of the LW detector were both thick and cold, they were very sensitive to the presence of cosmic rays, and slow to react to changes in fluxes. Therefore, for faint source detection with ISOCAM, it is necessary to discriminate non-Gaussian fluctuations of the signal from Gaussian ones, and to separate cosmic rays, i.e. glitches, from real sources. The method we developed for this purpose relies on the fact that these signal components, when measured by a given pixel, show different signatures in their temporal evolution, and can be identified using a multiscale transform, which separates the various frequencies in the signal. Once the "bad" components (i.e. glitches) are identified, they can be extracted from the temporal signal. The glitch-free signal can then be used to build the final image.

The detection of faint sources is then performed on this final image using again a wavelet transform of the signal, but this time spatially instead of temporally. We called this tool PRETI (Pattern REcognition Technique for Isocam data) because we use a temporal signature to recognize each signal component, which appears as a pattern in wavelet space.

In the first part of the paper, we describe the PRETI procedure. Then, we focus on the validation of this technique using Monte-Carlo simulations. These simulations were performed on a data set void of real sources in which we introduced simulated sources with random fluxes and positions, in order to estimate the following characteristics of an observation:

1. the sensitivity limit: the flux of the faintest detected source.
2. the photometric accuracy.
3. the completeness limit: the faintest flux for which all sources are detected or at least an established and important fraction of the total number of sources.
4. the false detection rate: the number of false sources due to glitches wrongly interpreted as sources, as a function of source strength.

A first set of simulations was already used in Aussel et al. (1999), based on this technique, but the used data set was not void of real sources, so that the fourth point above could not be addressed.

2. A brief discussion of the problem

The usual steps to analyze array data to construct a calibrated image are:

1. calibrating the data:
 - (a) extraction of the cosmic ray glitches by comparing successive readouts,
 - (b) subtraction of the signal due to dark currents,
 - (c) flat-fielding, by dividing the data by a flat-field extracted from the library,
 - (d) converting camera units into physical fluxes (Jy),
2. using a standard source detection algorithm, which allows us to estimate the background and the noise level, and fitting the Point Spread Function (PSF) to pixels showing a flux level higher than n times the noise standard deviation (rms).

As all ISOCAM surveys have been done using raster observations, we will not consider in this paper staring and CVF ISOCAM data (see Starck et al. 1999, for a general review of ISOCAM data calibration).

The simple calibration described above is successful when applied to bright objects (down to a few percent of the background level) but is inefficient when applied to faint source detection (below 1% of the background) with ISOCAM. At first order, this can be improved by modeling the flat-field, instead of using a library flat-field. The position of the lens of ISOCAM varies slightly between settings, and the optical flat-field varies as a function of the lens position by 2 to 20% from the center to the border of the array. In the case of empty fields (and more generally when most of the map covers an empty field), a simple median of the cube of data gives a very good flat-field, which allows us to reach a detection level of a few percent of the background level (Starck et al. 1999).

However, at second order, one encounters the main difficulty in dealing with ISOCAM faint source detection: the combination of the cosmic ray impacts (glitches) and the transient behavior of the detectors. For glitches producing single fast increases and decreases of the signal, a simple median filtering produces a fairly good deglitching. The ISOCAM glitch rate is one per second, and each glitch on average has an impact on eight pixels (Claret et al. 1999). However, 5 to 20% of the total number of readouts, depending on the integration time and the strength of the selection criterion, are affected by memory effects, which can produce false detections. Consequently, the main limitation here is not the detection limit of the instrument, which is quite low, but the false detections, whose number increases with the sensitivity.

Three types of glitches can be isolated, those creating:

1. a positive strong and short feature (lasting one readout only),
2. a positive tail (*fader*, lasting a few readouts),
3. a negative tail (*dipper*, lasting a several tens of readouts).

Figure 1 is a plot in camera units (ADU, for Analog to Digital Units) measured by a single pixel as a function of the number of readouts, i.e. time, which shows these three types of glitches: (a) three sharp “1” type, (b) a “fader” at readout 80 and lasting 20 readouts, (c) a “dipper” at readout 230 lasting 150 readouts.

The two first pixels are taken from a four by four raster observation of the Lockman hole, with a pixel field of view of 6 arc second, an individual integration time of 2.1 second, the LW3 filters (15 μm), a gain of 2, and 56 readouts for the first raster position and 27 readout for the others (observation number:03000102).

The last pixel is from another observation, with the same parameters except for the number of readouts per raster position, which is equal to 22 instead of 27 (observation number:02600404).

Finally, the signal measured by a single pixel as a function of time is the combination of memory effects, cosmic ray impacts and real sources: memory effects begin with the first readouts, since the detector faces a flux variation from an offset position to the target position (stabilization), then appear with long-lasting glitches and following the detection of real sources. Clearly one needs to separate all these components of the signal in each pixel before building a final raster map, and to keep the information of the associated noise before applying a source detection algorithm.

In Sect. 3, we will show that the concept of pattern recognition using a multi-resolution algorithm leads to an efficient calibration procedure, free of the major problems described above. Simulations and real data analysis will be presented in a Sect. 4.




Fig. 1. These three plots show a single detector response in Analog to Digital Units (ADU) as a function of time, expressed here in number of readouts, where each readout corresponds to 2 s. Three types glitches due to cosmic rays can be found here: **a)** top: the most common glitches, lasting only one readout. **b)** middle: a “fader”, occurring around readout 160. This glitch presents a slowly decreasing tail. It has been truncated above 100 ADUs, but its original intensity is 2558 ADUs. **c)** bottom: a “dipper”, beginning around readout 240. This glitch is followed by a memory effect lasting about 100 readouts. In these observations, the camera draws a mosaic on the sky (raster). Hence as long as an object, such as a star or a galaxy, is in the line of sight of a given pixel, the measured signal increases. A faint galaxy is visible on the bottom plot **c)** around readout 120, lasting only one position in the raster. Dashed lines mark different observation positions

3. “Pattern recognition”: A multi-scale approach

3.1. The “temporal detection technique” and its limitations

In Fig. 1c, at readout 120, a bump can be seen lasting one whole exposure time (30 readouts). This is the signal expected when the camera pointing is such that a source falls on the examined pixel. Thus we see that temporal detection of sources is a possible and useful alternative to the usual spatial detection. This temporal behavior of the observed flux by a pixel has the advantage of being dark and flat independent. Indeed, the flat and dark act as a multiplicative and an additive constant on the temporal signal, which does not affect the shape of the signal.

The redundancy inside the raster can be used for robust detection. For example, for a raster observation with successive array displacements corresponding to one half of the array on the sky, a source should be detected temporally four times (see Fig. 2). A criterion for a robust detection can be obtained by considering the number of times a source is detected. In the case of Fig. 2, the source can easily be detected by eye in the second and third plots while in the other two, the signal is too noisy.

We tested an automatic “temporal detection technique” which is described in Appendix A. Although its results are relatively robust, the technique still suffers from severe limitations:

1. *low signal to noise ratio (SNR)*: in a mosaic, a source generally extends over several pixels. Co-adding them allows us to increase the SNR. This is not possible in this technique.
2. *poor photometry*: because of the previous point and also due to the difficulty of estimating the background level.
3. *sources are split*: the signal of weak sources extended over several pixels (either because they are intrinsically extended or because of the Point Spread Function, PSF, or because of the camera distortion) is split, resulting again in a decrease of the SNR of the source.
4. *false detections*: false detections are possible for a redundancy of 10 or more, when searching for extremely faint objects, due to the large number of cosmic ray impacts.

In order to avoid these difficulties, it is necessary to identify the glitches with memory effects (faders and dippers), and extract them from the data, if possible without losing the associated information. Co-addition then becomes possible and a standard spatial source detection algorithm can be used, keeping in mind that noise is not homogeneously distributed on the map. This is exactly what PRETI allows us to do. New methods based on wavelet transforms have recently been developed for source extraction in an image (Bijaoui & Rué 1995), and successfully adapted for spectral analysis (Starck et al. 1997). Using such an approach, a temporal signal can be decomposed in its different components, selected from their frequency.

In terms of noise, in the temporal technique the noise standard deviation is divided by $\sqrt{N_r}$ (where N_r is the number of readouts per raster position), while it is divided by $\sqrt{N_r} * \sqrt{N_d}$ (where N_d is the number of redundancies inside the raster, i.e. the number of pixels which see the same sky position) for co-added data.

3.2. The Multi-Scale Vision Model

In the Multi-Scale Vision Model (Bijaoui & Rué 1995), an object in a signal is defined as a set of structures detected in the wavelet space. The wavelet transform algorithm used for such a decomposition is the so-called “à trous” algorithm, which allows us to represent a signal $D(t)$ by a simple sum of its wavelet coefficients w_j and a smoothed version of the signal c_p (see Appendix B for more details about the “à trous” algorithm)

$$D(t) = c_p(t) + \sum_{j=1}^p w_j(t). \quad (1)$$

The algorithm produces $p + 1$ arrays of the same size, each one containing only information at a given frequency




Fig. 2. Four pixel histories (top left pixel (4,12), top right pixel (12,12), bottom left pixel (20,12), and bottom right pixel (27,12)). The signal in ADU is plotted versus the readout number. The dotted lines indicate the change of camera pointing (i.e. between two dotted lines, the observed flux is the same). The same source is seen by the four pixels respectively at position 2,3,4,5. It can be relatively easily detected in pixel (12,12) and (20,12), while it is more difficult to see it in pixels (4,12) and (27,12)

band. In such signals, we define a “structure” as a group of connected significant (above a given threshold) wavelet coefficients. A complete description of how to estimate the significance of a wavelet coefficient, which depends on the nature of the noise, can be found in Starck et al. (1998). An object is described as a hierarchical set of structures. Two structures in a single object can be connected by “interscale-relation”. Consider two structures at two successive scales, S_j^1 and S_{j+1}^2 . Each structure is located on one of the individual arrays of the decomposition and corresponds to a region where the signal is significant. S_j^1 is said to be connected to S_{j+1}^2 if S_{j+1}^2 contains the pixel position of the maximum wavelet coefficient value of S_j^1 (i.e. the maximum position of the structure S_j^1 must also be contained in the structure S_{j+1}^2). A set of connected

structures is called an object in the interscale connectivity graph.

Once an object is detected in wavelet space, it can be isolated by searching for the simplest function which presents the same signal in wavelet space. The problem of reconstruction (Bijaoui & Rué 1995) consists then in searching for a signal V such that its wavelet coefficients are the same as those of the detected structure. By noting \mathcal{T} , the wavelet transform operator, and P_b , the projection operator in the subspace of the detected coefficients (i.e. P_b set all coefficients at scales and positions where nothing was detected to zero), the solution can be found by minimizing the following expression:

$$J(V) = \| W - (P_b \circ \mathcal{T})V \| \quad (2)$$




Fig. 3. The Multiscale Vision Model: contiguous significant wavelet coefficients form a structure, and following an inter-scale relation, a set of structures form an object. Two structures S_j, S_{j+1} at two successive scales belong to the same object if the position pixel of the maximum wavelet coefficient value of S_j is included in S_{j+1}

where W represents the detected wavelet coefficients of the signal. A complete description of algorithms for minimization of such a functional can be found in (Bijaoui & Rué 1995). In two dimensions, the method is identical. Figure 3 shows how several structures in different scales are linked together and form objects.

In ISOCAM data it is the cosmic ray glitches with memory effects which are the typical objects to be extracted from the time sequence of the signal. Indeed, the signal associated with a fader or a dipper type glitch is significant at several frequencies: a strong and rapid peak makes them significant in the highest frequency wavelet coefficient decomposition of the initial signal while the memory effect makes them significant in the lower frequency wavelet coefficients. Hence, the multi-scale approach is an ideal tool to discriminate glitches from real signal. We will call pattern recognition, the action of searching for objects showing properties typical of those expected for faders and dippers.

3.3. Pattern REcognition Technique for Isocam

The idea developed here is to use the multi-scale vision modeling for a decomposition of a signal into its main components. In practice, a simple object reconstruction from the detected structure in the wavelet space, as proposed in Bijaoui & Rué (1995), would produce poor results because of the strong confusion between the numerous objects that can be found in the data. Moreover, wavelet

transforms present a drawback: the wings of the wavelet function are negative (so that the integral of the function is zero) which implies that when a positive signal falls onto one wing of the wavelet function it produces a negative signal in the wavelet transform. The quality of the object reconstruction is good only when additional constraints are introduced, such as positivity constraint for positive objects, and negativity constraint for negative objects. An object is defined as positive (or negative) when the wavelet coefficient of the object, which has the maximum absolute value, is positive (or negative).

The problem of confusion between numerous objects can be solved when including a selection criterion in the detection of these objects. Using the knowledge we have about the objects, in this case, glitches, the problem of unknown object reconstruction is reduced to a pattern recognition problem, where the pattern is the glitch itself. We only search for objects which satisfy a given set of conditions in the Multi-Scale Vision Model (MVM). For example, finding glitches of the first type is equivalent to finding objects which are positive, strong, and with a duration shorter than those of the sources. The method that we use for the decomposition of the signal of a given ISOCAM pixel, $D(t_0... t_n)$, is summarized below:

1. detection of the glitches of the first type (i.e. few read-out glitches) in wavelet space: the corresponding signal, $C_1(t_0... t_n)$, is then subtracted from the initial data, D : $D_1 = D - C_1$. This is the first deglitching step.
2. detection of the negative components due to dippers: the multi-scale vision model is applied to D_1 , hence negative objects are detected and the reconstructed signal, $C_2(t_0... t_n)$, is subtracted to the output of the previous step: $D_2 = D_1 - C_2$. This is the second deglitching step where throughs following glitches are corrected.
3. detection of the positive components due to faders and dippers: this step must be done carefully, since sources also produce positive components in the signal. Only positive objects lasting much longer or much less than the time spent on a given position on the sky are automatically considered as glitches. The output signal, $C_3(t_0... t_n)$, is then subtracted again from the previous signal: $D_3 = D_2 - C_3$.
4. detection of a very strong positive signal on scales where sources are expected. This step is done in preparation for the baseline subtraction; the final source detection is not done at this stage. The multiscale vision model is applied to D_3 and strong positive objects with a correct temporal size are reconstructed: we obtain $C_4(t_0... t_n)$, and we calculate $D_4 = D_3 - C_4$.
5. baseline subtraction: the signal D_4 contains only noise and temporally undetectable faint sources. The baseline is easily obtained by convolving D_4 by a low frequency pass band filter. We obtain $C_5(t_0... t_n)$.




Fig. 4. These two plots show the signal of a single pixel as a function of time before calibration (top, flux in ADU) and after calibration (bottom, flux in ADU/gain/second). The trough following the second glitch (dipper) has disappeared and the remaining signal contains only Gaussian noise (photon noise + readout noise) plus sources (one relatively bright source is located at readout 120, fainter sources will only appear after co-addition of all pixels having seen the same sky position in the final map)

6. The residual noise is obtained by $C_6 = D_4 - C_5$; its mean value is zero.

Finally, the set $(C_1, C_2, C_3, C_4, C_5, C_6)$, represents the decomposition of the signal into its main components. Note also that the input signal D is equal to the sum of all components:

$$D = \sum_{i=1}^6 C_i. \quad (3)$$

A deglitched signal is obtained by:

$$D_g = D - C_1 - C_2 - C_3. \quad (4)$$

For faint source detection, we use the signal $D_b = C_4 + C_6$, which is background, dark, and glitch free. The background has been subtracted, and glitches with their long duration effects have been suppressed. Applying the pattern recognition method to all detector pixels, we obtain a cube $D_b(x, y, t)$. All other components are kept in the cubes C_i . The baseline suppression presents several advantages: first, the final raster map is dark-corrected without the need of a library dark, since we end up with a mean zero level for each pixel. This is particularly important when the library dark is not good enough, and induces visual artifacts (Starck et al. 1999). Second, the flat-field accuracy only affects the photometry of the sources but not the background level, which is extracted in the baseline. Thus, its influence in the final calibration is decreased.

3.4. Example

Figure 4 (bottom) presents the result obtained with this method. The decomposition of the original signal (Fig. 4

top) into its main components is shown in Fig. 5: (a), (b), and (d) are features subtracted from the original signal (short glitches, dipper, and baseline, respectively), which present no direct interest for faint source detection, and only (c) and (e) (bright sources and noise plus faint sources, respectively) are kept for building the final image. The noise must also be kept because faint sources are often detectable only after co-addition of the data. The simple sum of the five components is exactly equal to the original data (see Fig. 4 top). The calibrated background free data (see Fig. 4 bottom) are then obtained by addition of (c) and (e).

3.5. Transient correction

Three kinds of transients must be distinguished:

1. a long term transient at the beginning of the observation. It can be either a downward or an upward transient depending on the flux level of the previous observation. If the difference between the present background level (which dominates the signal in faint source observations) and the previous background level is high, then the transient can affect several hundred frames. Long term transients have no effect on source detection when using the multi-resolution approach, because they are eliminated with the baseline.
2. an upward transient each time a pixel points in the direction of a source. There is presently no physical model which describes the ISOCAM upward transient. This type of transient affects mainly the photometry. Objects with a flux at the theoretical detection limit are not detected because the signal measured is only a fraction (typically 60%) of the signal after stabilization (Starck et al. 1999).
3. a downward transient after each source, which can produce ghosts when following bright sources, since the downward transient may remain above the noise level even after the change of camera pointing.

For an automatic source detection method, the last type of transient must be corrected for. Physical models exist for downward transients and several methods may be used (Starck et al. 1999). In our case, a very trivial approach can also be used, which consists of treating the reconstructed temporal objects. Indeed, we can assume that the part of the object which appears after the displacement of the array is the transient, and it can be eliminated by a simple thresholding.

Figure 6 bottom shows the result after such a treatment. A signal containing a source (top) shows a strong downward transient, which is very clear after deglitching and baseline subtraction (middle). The three successive positions on the array are affected by the transient, which induces ghosts in the final image if they are not




Fig. 5. Decomposition of the signal into its main components: **a)** short glitch, **b)** trough of a dipper, **c)** bright source, **d)** baseline, **e)** noise plus faint sources. The simple sum of the five components is exactly equal to the original data (see Fig. 2). The calibrated background free data are obtained by addition of signals **c)** and **e)**. Figure **c)** shows the reconstruction of a source approximated as a Gaussian, but sources are kept in the signal and their shape differ from one source to the other




Fig. 6. Examples of transient corrections. Top: signal containing a strong source. Middle: signal after deglitching and subtraction of the baseline. The dashed line shows the configuration limits and crosses indicate the readouts which have been masked. Bottom: signal after the transient correction

removed. The correct shape of the source (bottom) is obtained through the technique developed by Abergel et al. (1996).

3.6. The Multiscale Median Transform

The presented method produces good results but requires a long computation time. A similar but faster method, producing results of equivalent quality and avoiding the delicate problem of the negative wings of wavelet functions, is to use the Multi-Resolution Median Transform (MMT) (Starck et al. 1998) instead of the wavelet transform. No confusion between positive and negative objects is possible because this multi-resolution transform does not present the ringing drawback. The MMT has been proposed for data compression (Starck et al. 1996), and it has also recently been used for ISOCAM short glitch suppression (Starck et al. 1999).

The MMT algorithm is relatively simple. Let $\text{med}(S, n)$ be the median transform of a one-dimensional signal S within a window of dimension n . If we define, for N_s resolution scales, the coefficients:

$$c_i = \begin{cases} S & \text{if } i = 1 \\ \text{med}(S, 2^{i-1} + 1) & \text{if } i = 2, N_s \end{cases} \quad (5)$$

$$w_i = c_{i-1} - c_i \quad \text{for } i = 2, N_s \quad (6)$$

we can expand the original signal similarly to the “à trous algorithm”:

$$S = c_p + \sum_{i=2, N_s} w_i, \quad (7)$$

where c_p is the residual signal.

Applying the multi-scale vision model (MVM) with the MMT, the object reconstruction is straightforward, because no iteration is needed to get a good quality reconstruction.

Figure 7 shows the comparison between the “à trous” wavelet transform (on the left) and the MMT (on the right) of the input signal in Fig. 1c. In these diagrams, the scale is represented as a function of the time. We can see that with the MMT, it is easier to distinguish sources from glitches. Figure 8 shows the result of MVM using the MMT instead of the “à trous” algorithm. From (a) to (e), we see the reconstructed source, the dipper, the baseline, the sum of the glitches and the baseline (i.e. non-interesting data), and the original signal from which the signal (d) has been subtracted.

3.7. Conclusion

Once the calibration is done, the raster map can normally be created, with flat field correction, and all data co-added. The associated rms map can now be used for the detection, which was impossible before due to the strong effect of residual glitches. Since the background has been removed, simple source detection can be done just by comparing the flux in the raster map to the rms map.

4. Source detection

Once all data have been calibrated, the final raster map $R(x, y)$ and its associated rms map $R_\sigma(x, y)$ can be created. If several rasters of the same field are available, they can be co-added in order to improve the signal to noise ratio. The noise in $R(x, y)$ (i.e. $R_\sigma(x, y)$) is non-homogeneously distributed over the map, first because some pixels have been masked (short glitches) and second because some areas of the field (particularly the border of the mosaic) present low redundancy (few readouts per sky position). For this reason the noise around the border of the image can be relatively high with respect to the noise toward the image center. Therefore, if we made the simple hypothesis of uniform noise (for instance Gaussian noise of standard deviation σ), it would lead to a large amount of false detections on the border. The correct solution is to use the $R_\sigma(x, y)$ map. In order to detect faint sources on the final image, we can use the multi-scale vision model (MVM) in two-dimensions. This time we use the “à trous” algorithm because the linearity of the wavelet transform allows us to derive a robust modelling of the noise in wavelet space (using the rms map $R_\sigma(x, y)$), which is impossible using the MMT. Moreover, in this case the artefacts around sources are negligible since we have no strong sources.

For each wavelet coefficient $w_j(x, y)$ of R , the exact standard deviation $\sigma_j(x, y)$ has to be calculated from the root mean square map $R_\sigma(x, y)$.

A wavelet coefficient $w_j(x, y)$ is obtained by the correlation product between the image R and a function g_j :

$$w_j(x, y) = \sum_k \sum_l R(x, y) g_j(x + k, y + l) \quad (8)$$

then we have:

$$\sigma_j^2(x, y) = \sum_k \sum_l R_\sigma^2(x, y) g_j^2(x + k, y + l). \quad (9)$$

In the case of the “à trous” algorithm, the coefficients $g_j(x, y)$ are not known exactly, but they can be computed by taking the wavelet transform of a Dirac (w^δ , in our notation):

$$g_j(x, y) = w_j^\delta(x, y). \quad (10)$$

Then the map σ_j^2 is calculated by correlating the square of the wavelet scale j of w^δ by $R_\sigma^2(x, y)$. A wavelet coefficient is significant if:

$$|w_j(x, y)| > N_\sigma \sigma_j(x, y). \quad (11)$$

N_σ is a parameter fixing the confidence level (generally taken equal to 3). Once this step is performed, the object selection and their reconstructions can be done as described in Bijaoui & Rué (1995). One can therefore produce a map containing only the reconstructed objects, i.e. the sources (galaxies, stars) that we were looking for. This image can be used for comparison at other wavelengths.

Finally, the outputs of PRETI are numerous and contain all information at all scales divided into several cubes of data and images in FITS format, but the most commonly used outputs are the following:

1. the final image and its associated rms image.
2. the image of the reconstructed objects
3. the list of objects with their position, flux and flux error (assuming no transient error), the sigma level of the detection and the scale at which the object was detected (size of the object).

The confidence level associated with the faint sources detected with PRETI cannot be directly understood as a usual signal-to-noise ratio typical of a Gaussian noise. In fact, because of the presence of residual glitches, the detection level is not fixed by the overall rms but by the level at which false detections begin to appear. One can check the robustness of the source detection by using a high detection level and comparing the source list with the brightest sources of the corresponding optical image. However, to tackle adequately the faint source detection, the only solution is to tune the method with simulations.

A field where this faint source detection method is widely applied is the study of source number counts in galaxy surveys, as for example the ISOCAM survey in the Hubble Deep Field region (Aussel et al. 1999). We will explain in the next section how simulations allows us to determine the accuracy of the so called $\log N - \log S$ diagrams, which show the number N of sources as a function of the emitted flux S .

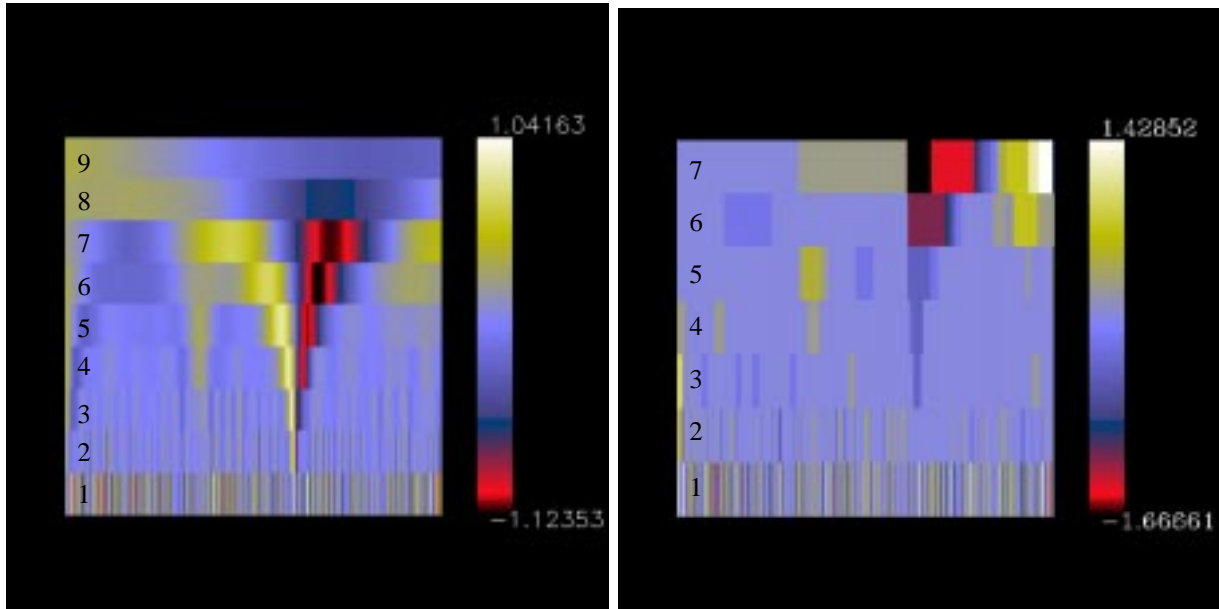


Fig. 7. Comparison between the “à trous” wavelet transform (left) and the multiresolution median transform (right) of the signal of Fig. 1c. Resolution scale is represented versus the time. We note that the separation between the source and the glitch is improved using the MMT




Fig. 8. Decomposition of the signal into its main components using the MMT: **a)** source, **b)** dipper, **c)** baseline, **d)** Sum of the glitches and the baseline (i.e. non interesting data), and **e)** original signal minus **d)**

5. Simulations

Simulations are essential to tune the method in order to detect the faintest sources without excessive false detections. They allow one to compute the following quantities:

1. *sensitivity limit*: the flux of the faintest source detected;
2. *photometric accuracy*: the response of the detectors is generally not stabilized when pointing towards faint

sources because i stabilization is very slow. Hence the uncertainty on the photometry depends on the quality of the algorithm for transient correction (Abergel et al. 1996).

3. *completeness limit*: the faintest flux for which all sources or at least a known fraction of sources are detected.




Fig. 9. Simulation of the ISO-HDF mosaic (I): (left) this image contains only simulated photon noise plus readout noise, i.e. Gaussian noise, (right) image of the ISO-HDF mosaic without any noise and with simulated sources according to a distribution without evolution (Franceschini 1997). Sources ranges from $0.1 \mu\text{Jy}$ to 1 mJy




Fig. 10. Simulation of the ISO-HDF mosaic (II): (left) sum of the two previous images, i.e. simulated sources plus readout and photon noise, (right) image of the ISO-HDF mosaic simulated from a staring observation, i.e. all sources of noise are present but no source at all is present (this image can be used to estimate the number of false detections due to glitches, since it does not contain any real sources)

4. *rate of false detections*: the number of false detections due to glitches mimicking sources.

We describe below the details of the simulations.

5.1. Noise and cosmic rays simulation

If there are no sources, the temporal signal of the pixel contains only instrumental noise and cosmic ray effects. Due to the lack of accuracy of our interpretations of the physical effects of cosmic rays on the detector, we are compelled to use real ISOCAM data. A first possibility is to add simulated sources to the real observation: this allows us to test and calibrate the photometry and verify the completeness at a given flux by comparing the number of new detections with the number of introduced simulated sources. This method was applied in the study of the ISOCAM image of the Hubble Deep Field (Aussel et al. 1999). However, this technique does not allow us to measure the false detection rate and to test the possibility of recovering the $\log N - \log S$ relation if it is sensitive to confusion. Thus we needed a long observation void of sources, where the only sources are the simulated sources intentionally added to the dataset. This very long pointed observation (“staring”) has been cut in order to create a false “raster” observation. If a pixel detects a source, its flux remains above the background level throughout the observation, and this has no effect on the source detection algorithm since the low frequency signal is subtracted. The “staring” observation has to be at least as long as the raster to be simulated. Moreover, the rate of cosmic rays must be compatible with that of the observation and the individual integration time must be the same, since the behavior of the detectors depends on it. Finally, the high frequency photon noise level of the two observations must be compatible. Such an observation was acquired for us during the calibration of ISOCAM toward the end of the ISO mission, with both the LW2 and LW3 filters.

5.2. Source simulation

The behavior of a source on a detector is a step function i.e. the signal of a pixel increases when observing the source and afterwards decreases down to the background level. The presence of transients modifies this behavior. We used the Abergel et al. (1996) inverse transient model to simulate the sources seen by ISOCAM. To take into account the PSF effect, i.e. the distribution of the flux of a point source among the nearby pixels, we adopt the (Okumura 1997) model. In this way we are able to simulate a source once its position on the detector is known. An observation is therefore composed of a set of sources whose positions follow a uniform probability density function.

To study the completeness limit and the photometric accuracy, we generate a list of sources with uniform

flux. Several lists must be generated to ensure that peculiar source positions (e.g. a pixel affected by numerous glitches) do not affect the result. Fake observation data are then created for several flux levels.

To test the validity of the number counts ($\log N - \log S$ or dN/dS) obtained from one observation, we can analyze simulated observations which contain random lists of fake sources whose fluxes follow the theoretical $\log N - \log S$. In Sect. 6, we apply this method to the case of the Hubble Deep Field, North.

5.3. Specificity of a simulation

Ideally, to reach the ultimate limit of the instrument, a new set of simulations should be produced and analyzed for each observation, since the results depend on the parameters of the observations as well as on the background level and glitch rate. However, typical cases can be analyzed and used as templates for other observations. We have performed detailed simulations (one hundred simulations per observation) for two template cases: the “ISO-HDF”, which corresponds to what we call ultra-deep observations with a very large redundancy (see below), and the “Deep Survey” (Elbaz et al. 1998), for shallower observation with less redundancy and spatial resolution.

6. The case of the Hubble Deep Field North

The Hubble Deep Field North was observed by ISO on July 1996 (Serjeant et al. 1997) and it was analyzed by several independent groups (Goldschmidt et al. 1997; Désert et al. 1999; Aussel et al. 1999).

This work completes a previous paper on the ISO-HDF North (Aussel et al. 1999) with new simulations on an ideal dataset at $15 \mu\text{m}$ (filter LW3, $12 - 18 \mu\text{m}$, see Figs. 9, 10, 11). In order to check if the conditions were similar for both observations (the staring observation and the real mosaic on the HDF field), we first performed an analysis of the staring observation in order to determine the percentage of readouts affected by glitches of type 2 and 3 (faders and dippers). In the case of the ISO-HDF (see Table 1), the percentage of data affected because of faders and dippers is about 20%, close to the fraction of pixels lost because of glitches of type 1. In the case of Deep Survey-like observations, for comparison where the redundancy per sky position is much lower (about 3 to 6, instead of 64), one cannot set such strong criteria for the correction of glitches with memory effects and the typical fraction of corrected pixels is of the order of 5% (the fraction of lost pixels, however remains identical). Finally, we also checked that the Gaussian plus readout noise mean is comparable in both observations, which is indeed the case when considering the same integration time as shown in Table 1.




Fig. 11. ISO-HDF mosaic: (left) full simulated ISO-HDF image using the staring observation plus simulated sources following a distribution without evolution (Franceschini 1997), i.e. sum of Fig. 8 (right) plus Fig. 7 (right), (right) real image of the ISO-HDF. We find more sources in the real image than in the simulated image, indicating strong evolution

Table 1. Comparison of the noise level and rate of cosmic ray glitches during the ISO-HDF observation and during a staring observation used for building a simulation of the same ISO-HDF mosaic in the LW3 band ($12 - 18 \mu\text{m}$)

Observation	ISO-HDF	Simulation
Noise level (ADU/G/s)	0.229	0.232
Masked pixels	19.4%	19.1%
Corrected pixels:	19.3%	18.9%
Faders	5.2%	5.5%
Dippers	14.1%	13.4%

In order to quantify the effect of incompleteness plus photometric uncertainty on the number counts, we built several ISO-HDF simulated images including simulated sources whose flux distribution followed that proposed in Franceschini et al. (1997). We stress here that this does not influence the output number counts, but on the contrary allows us to check if after applying PRETI, the slope of the number counts was close to the one used in the input. The main uncertainty here comes from the accuracy reached in the photometry of the sources, which redistributes the sources in each flux bin.

We use a lower limit of $S_l = 0.1 \mu\text{Jy}$, which is much lower than the sensitivity of our observations, and an upper limit $S_u = 1 \text{ mJy}$, for the fluxes of the fake sources. We then simulate several fake mosaics with fluxes distributed as described above. Fluxes in Jy are converted into ISOCAM units (ADU/gain/second) following the standard conversion table from the ISO cookbook (ISO-Team, 1994) ($1 \text{ ADU/g/s} = 1.96 \text{ mJy}$ with LW3).

We project each source on the detector for each pointing of the camera by taking into account the field distortion and the point spread function (using the model of Okumura 1997). This allows us to build a cube of images, i.e. an image of 32×32 pixels for each pointing of the satellite, that we then multiply by the flat-field computed from the original data without simulated sources. Finally, we add to this cube the mean background level of the staring observation in order to model the transient behavior of the sources, which depends on the total flux level of each detector, using the model of Abergel (1996). Finally, we subtract again the mean background level to this cube of simulated sources in order to add it to the real dataset of the staring observation, which contains the real background with noise and cosmic rays.

Figure 12 shows the fraction of false detections as a function of flux limit using a detection threshold of $7\tau_w$ and $5\tau_w$, where τ_w is the noise level in wavelet space. For each simulation, we have built a main and a supplementary list of sources as mentioned in the paper and found that in both cases the rate of false detection down to the completeness limit is only 2%.

Hence in the main list of sources extracted from the ISO-HDF (21 sources), which was built using the $7\tau_w$ threshold, the completeness limit is $200 \mu\text{Jy}$ while the sensitivity limit is $50 \mu\text{Jy}$ with a rate of false detection close to 2%. But in the supplementary list (a total of 46 sources including the previous list), which goes down to $5\tau_w$, the completeness limit is $100 \mu\text{Jy}$ with about the same rate of false detection of 2%, which we could not measure previously. Hence, we can now merge these two source lists




Fig. 12. Fraction of false detections as function of the flux limit of the sample, for a detection threshold of $7 \tau_w$ (lower) and $5 \tau_w$ (upper)

into one single list of 46 sources, the completeness limit of which is $100 \mu\text{Jy}$ instead of $200 \mu\text{Jy}$, and contains statistically only one false source (2% of 46 sources) with a confidence level of 95%. At fainter fluxes, the number of false detections increases very rapidly in the supplementary list. For a limit of $20 \mu\text{Jy}$ it varies between 5 and 30% according to the simulations. On a list of one hundred sources this implies 10 false detections between 20 and $80 \mu\text{Jy}$.

7. Conclusion

We have developed a tool for faint sources detection in ISOCAM data, which proved to be particularly well adapted for the detection of sources at the few tens of μJy level in the presence of glitches with memory effects. We created simulated datasets in order to test the robustness of the technique and found a quantitative way to estimate the quality of the source lists extracted from deep surveys, while the signal-to-noise ratio alone would be misleading. We applied this technique to a simulation of the ISO-HDF North and found that the completeness limit at $15 \mu\text{m}$ is $100 \mu\text{Jy}$ with 2% of false detections due to remnant glitches. New results already obtained with PRETI will follow this paper in the near future.

Appendix A: Temporal method for faint source detection

The principle

As glitches are the main limitations for faint ISOCAM source detection, not the noise, it is clear that an analysis

on the final raster image with a standard method (detection at $k\sigma$ the noise level + background) would lead to poor results if the glitches with transients have not been removed. In order to avoid other problems (dark current subtraction, flat field, transient and long drift correction, etc.), a solution is to perform a temporal source detection technique rather than a standard source detection technique on the final raster map. The temporal source detection method is based on the fact that the flux observed by a single detector increases when the detector points toward a source and decreases when the camera is moved to the subsequent phase of a raster observation. This temporal behavior of the flux observed by a detector has the advantage of being dark current and flat field independent. Indeed, the flat field and dark current act as a multiplicative and an additive constant on the total temporal signal, and do not effect the shape of the signal. Thus, the signature of a source can be identified.

Temporal detection

Short glitches (i.e. first type) can be easily removed by masking the position where they appear (Starck et al. 1999). For each pixel (x, y) , we indicate the deglitched data as $D(x, y, c, r)$ and the corresponding mask as $M(x, y, c, r)$ (0 if the position is masked, 1 otherwise), where c and r indicate respectively the configuration (raster position) and the readout number in this configuration. Values corresponding to the same sky position and the same configuration are averaged:

$$I(x, y, c) = \frac{\sum_r M(x, y, c, r) D(x, y, c, r)}{\sum_r M(x, y, c, r)}. \quad (12)$$

The temporal noise $\sigma(x, y)$ is estimated for each pixel independently using a k -sigma clipping method, so the noise on the mean value of signal in a configuration is given by:

$$\sigma_I(x, y, c) = \frac{\sigma(x, y)}{\sqrt{\sum_r M(x, y, c, r)}}. \quad (13)$$

The detection is done by calculating the signal:

$$W(x, y, c) = I(x, y, c) - \frac{1}{2}(I(x, y, c-1) + I(x, y, c+1)) \quad (14)$$

and its associated noise:

$$\sigma_W(x, y, c) = \sqrt{(\sigma_I(x, y, c))^2 + \left(\frac{1}{2}\sigma_I(x, y, c-1)\right)^2 + \left(\frac{1}{2}\sigma_I(x, y, c+1)\right)^2}. \quad (15)$$

Then we consider we have a detection at pixel (x, y) and at the configuration c if:

$$W(x, y, c) > k\sigma_W(x, y, c), \quad (16)$$

where, in general, k is taken equal to 3. If a source is detected at position (x, y, c) we put $I(x, y, c) = 1$, otherwise $I(x, y, c) = 0$. We can therefore coadd the C $I(x, y, c)$ matrixes in order to obtain a matrix of detections with size equal to that of the total image: $Image(\xi, \eta)$ indicates how many times a source has been detected at the sky position (ξ, η) during the raster observation. For instance, for a raster observation with half overlapping, $Image(\xi, \eta)$ can take the integer values between 0 and 4.

Constraints for a robust detection

The detection has been made under the assumption of Gaussian noise. Due the large number of glitches, false detections will occur. Two parameters can be adjusted in order to limit the number of false detections:

1. the detection level (k parameter). By default, the detection is done at 3σ (it corresponds to a false detection probability of 0.25%). Increasing the detection level eliminates false detections (but also weaker objects).
2. the number of required redundancy N_r . For a raster made by overlapping half the array, a source should be detected four times (two times if it is on the border of the raster image). Fixing a minimum of two detections should suppress most of the false detections.

A robust detection is performed by comparing $Image(x, y)$ to the number N_r of required redundancy. A high redundancy allows one to increase N_r and improve the robustness of the detection. We point out that $Image(x, y)$ is independent of the background level, the flat field, and the dark.

Conclusion

Once (Goldschmidt et al. 1997; Serjeant et al. 1997) the detection is done, sources must be extracted with astrometric and photometric information. This is done using the final calibrated raster image (see Siebenmorgen et al. 1996, for a complete description of each calibration step). For observations with the six arc second lens, the PSF is mainly contained in one single pixel. So PSF fitting does not help, and the flux of an object can be obtained by integrating the flux in a small box around the detected position, using an estimate of the background. The gain variation due to dippers and faders will have an effect on the accuracy of the photometry, because it modifies the background on a individual pixel. To summarize this approach, the advantages are that the detection is relatively robust and independent of the dark current and the flat field, while the drawbacks are:

1. the photometry is poor;
2. the temporal detection does not allow the use of correlation between adjacent pixels, which is needed for extended weak sources detection;
3. data cannot be coadded before detection.

To overcome these problems, the only way is to correct the data from the gain variation due to faders and dippers.

Appendix B: The “À Trous” wavelet transform algorithm

In a wavelet transform, a series of transformations of an image is generated, providing a resolution-related set of “views” of the image. The properties satisfied by a wavelet

transform, and in particular the *à trous* wavelet transform (“with holes”, so called because of the interlaced convolution used in successive levels: see step 2 of the algorithm below) are further discussed in (Starck et al. 1998).

We consider sampled data, $\{c_0(k)\}$, defined as the scalar product at pixels k of the function $f(x)$ with a scaling function $\phi(x)$ which corresponds to a low pass band filter:

$$c_0(k) = \langle f(x), \phi(x - k) \rangle. \quad (17)$$

The scaling function is chosen to satisfy the dilation equation:

$$\frac{1}{2}\phi\left(\frac{x}{2}\right) = \sum_k h(k)\phi(x - k) \quad (18)$$

h is a discrete low pass filter associated with the scaling function ϕ . This means that a low-pass filtering of the image is, by definition, closely linked to another resolution level of the image.

The smoothed data $c_j(k)$ at a given resolution j and at a position k is the scalar product

$$c_j(k) = \frac{1}{2^j} \langle f(x), \phi\left(\frac{x - k}{2^j}\right) \rangle. \quad (19)$$

This is consequently obtained by the convolution:

$$c_j(k) = \sum_l h(l) c_{j-1}(k + 2^{j-1}l). \quad (20)$$

The signal difference w_j between two consecutive resolutions is:

$$w_j(k) = c_{j-1}(k) - c_j(k) \quad (21)$$

or:

$$w_j(k) = \frac{1}{2^j} \langle f(x), \psi\left(\frac{x - k}{2^j}\right) \rangle. \quad (22)$$

Here, the wavelet function ψ is defined by:

$$\frac{1}{2}\psi\left(\frac{x}{2}\right) = \phi(x) - \frac{1}{2}\phi\left(\frac{x}{2}\right). \quad (23)$$

For the scaling function, $\phi(x)$, the B-spline of degree 3 was used in our calculations. We have derived a simple algorithm in order to compute the associated wavelet transform:

1. We initialize j to 0 and we start with the data $c_j(k)$.
2. We increment j , and we carry out a discrete convolution of the data $c_{j-1}(k)$ using the filter h . The distance between the central pixel and the adjacent ones is 2^{j-1} .
3. After this smoothing, we obtain the discrete wavelet transform from the difference $c_{j-1}(k) - c_j(k)$.
4. If j is less than the number p of resolutions we want to compute, then go to step 2.
5. The set $\mathcal{W} = \{w_1, \dots, w_p, c_p\}$ represents the wavelet transform of the data.

The most general way to handle the boundaries is to consider that $c(k + N) = c(N - k)$. But other methods can be used such as periodicity ($c(k + N) = c(k)$), or continuity ($c(k + N) = c(N)$). Choosing one of these methods

has little influence on our general restoration strategy. We used continuity.

A series expansion of the original signal, c_0 , in terms of the wavelet coefficients is now given as follows. The final smoothed array $c_p(x)$ is added to all the differences w_j :

$$c_0(k) = c_p(k) + \sum_{j=1}^p w_j(k). \quad (24)$$

This equation provides a reconstruction formula for the original signal.

At each scale j , we obtain a set $\{w_j\}$. This has the same number of pixels as the input signal.

The above *à trous* algorithm has been discussed in terms of a single index, x , but is easily extendable to two-dimensional space. The use of the B_3 spline leads to a convolution with a mask of 5×5 :

$$\begin{pmatrix} \frac{1}{256} & \frac{1}{64} & \frac{3}{128} & \frac{1}{64} & \frac{1}{256} \\ \frac{1}{64} & \frac{1}{16} & \frac{3}{32} & \frac{1}{16} & \frac{1}{64} \\ \frac{3}{128} & \frac{3}{32} & \frac{9}{64} & \frac{3}{32} & \frac{3}{128} \\ \frac{1}{64} & \frac{1}{16} & \frac{3}{32} & \frac{1}{16} & \frac{1}{64} \\ \frac{1}{256} & \frac{1}{64} & \frac{3}{128} & \frac{1}{64} & \frac{1}{256} \end{pmatrix} \quad (25)$$

In one dimension, this mask is: $(\frac{1}{16}, \frac{1}{4}, \frac{3}{8}, \frac{1}{4}, \frac{1}{16})$.

To facilitate computation, a simplification of this wavelet is to assume separability in the 2-dimensional case. In the case of the B_3 spline, this leads to a row by row convolution with $(\frac{1}{16}, \frac{1}{4}, \frac{3}{8}, \frac{1}{4}, \frac{1}{16})$; followed by column by column convolution.

As for the one dimensional case, an exact reconstruction is obtained by adding all the scales and the final smoothed array:

$$c_0(x, y) = c_p(x, y) + \sum_{j=1}^p w_j(x, y). \quad (26)$$

Acknowledgements. We wish to thank Suzanne Madden for her kind help in the revision of the manuscript. We also would like to thank the anonymous referee whose comments had helped us to improve the quality of the paper.

References

- Abergel A., Bernard J.-P., Boulanger F., 1996, A&A 315, L329
 Aussel H., Elbaz D., Cesarsky C., Starck J., 1999, A&A 342, 313
 Bijaoui A., Rué F., 1995, Sign. Proc. 46, 229
 Cesarsky C., Abergel A., Agnese P., et al., 1996, A&A 315, L32
 Claret A., Dzitko H., Engelmann J., Starck J.-L., 1999, Exper. Astron. (in press)
 Désert F., Puget J.-L.D., Clements M.P.A.A., Bernard J., Cesarsky C., 1999, A&A 134, 342
 Djorgovski S., Thompson D., 1992, in IAU Symposium 149: The stellar population of galaxies, Barbuy B., Renzini A. (eds.). Kluwer, Dordrecht, p. 337
 Elbaz D., Aussel H., Cesarsky C., et al., 1998, in Cox P., Kessler M. (eds.), ESA Conference: Universe as seen by ISO, ESA Special Publications series (SP-427), astro-ph 9902229
 Franceschini A., 1997, in Mamon G., Thuân T.X., Vãn J.T.T. (eds.), Extragalactic Astronomy in the Infrared. Éditions Frontières, pp. 509–519
 Goldschmidt P., Oliver S.J., Serjeant S.B.G., et al., 1997, MNRAS 289, 465
 Houck J.R., Schneider D.P., Danielson G.E., et al., 1985, ApJL 290, L5
 Houck J.R., Soifer B.T.G.G.N., et al., 1984, ApJL 278, L63
 ISO-Team, 1994, “ISOCAM Observer Manual”, http://isowww.estec.esa.nl/manuals/iso_cam/cam_om_1.html
 Kessler M.F., Steinz J.A., Anderegg M.E., et al., 1996, A&A 315, L27
 Okumura K., 1997, ISOCAM PSF Report, ESA/CAM IDT, http://isowww.estec.esa.nl:80/instr/CAL/cal_wksp
 Serjeant S.B.G., Eaton N., Oliver S.J., et al., 1997, MNRAS 289, 457
 Soifer B.T., Neugebauer G., Helou G., et al., 1984a, ApJL 283, L1
 Soifer B.T., Rowan-Robinson M., Houck J.R., et al., 1984b, ApJL 278, L71
 Starck J., Abergel A., Aussel H., et al., 1999, A&AS 134, 135
 Starck J., Murtagh F., Bijaoui A., 1998, Image Processing and Data Analysis: The Multiscale Approach. Cambridge University Press, Cambridge (GB)
 Starck J., Murtagh F., Pirenne B., Albrecht M., 1996, PASP 108, 446
 Starck J., Siebenmorgen R., Gredel R., 1997, ApJ 482, 1011
 Vigroux L., Charmandaris V., Gallais P., et al., 1998, in Cox P., Kessler M. (eds.). ESA Conference: Universe as seen by ISO, ESA Special Publications series (SP-427)