

HIRES AND BEYOND

395743

51-82
40921
1/10

J. W. FOWLER AND H. H. AUMANN

*Infrared Processing and Analysis Center, MS 100-22**California Institute of Technology, Jet Propulsion Laboratory**Pasadena, CA 91125*

ABSTRACT The High-Resolution image construction program (Hi-Res) used at IPAC is based on the Maximum Correlation Method. After HiRes intensity images are constructed from IRAS data, additional images are needed to aid in scientific interpretation. Some of the images that are available for this purpose show the fitting noise, estimates of the achieved resolution, and detector track maps. Two methods have been developed for creating color maps without discarding any more spatial information than absolutely necessary: the "cross-band simulation" and "prior-knowledge" methods. These maps are demonstrated using the survey observations of a 2×2 degree field centered on M31. Prior knowledge may also be used to achieve super-resolution and to suppress ringing around bright point sources observed against background emission. Tools to suppress noise spikes and for accelerating convergence are also described.

1. INTRODUCTION

The high-resolution image construction program used at IPAC is based on the Maximum Correlation Method (MCM; see Aumann, Fowler, and Melnyk, 1990). This program uses IRAS survey and additional observation data to construct images in the four IRAS survey wavelength bands. These images are intensity maps covering typically one to four square degrees. The effective beam size varies significantly over these images because of the manner in which the IRAS telescope gathered infrared data. This makes interpreting the images and obtaining color ratios especially challenging. Complications are added by the nature of the celestial background (the varying zodiacal background) and the detectors. The more important effects caused by the detectors and the orbital environment are: hysteresis, residuals of radiation hits, saturation for bright sources, detectors of different noise amplitude, and baseline drift due to $1/f$ noise (*IRAS Explanatory Supplement 1988*).

In this paper we discuss additional tools and techniques that have been developed at IPAC over the last two years to address these difficulties. Some of these are byproducts of the intensity imaging (e.g., correction-factor variance and coverage maps), while others have been added recently to the default products (e.g., beam-sample maps and detector track maps). Some techniques that have been developed are now employed by default (e.g., automatic flux-bias pro-

Fig. 4: Color-ratio map of M31
from cross-band simulation
(100 micron/60 micron)




Fig. 3: Detector track map for Fig. 1




Fig. 2: Correction-factor variance
map for Figure 1




Fig. 1: Intensity image of M31 at 100
microns after 20 iterations



cessing), but most remain available by special-request only (e.g., point-source masking, scan-direction masking, spike median filtering, n-sigma data editing, point-source ringing suppression, color ratio and multi-band intensity maps). Some techniques require special preparation by the investigator (e.g., simulation and prior-knowledge input).

2. CORRECTION-FACTOR VARIANCE MAPS

The MCM iteratively constructs an image that reproduces the observed data to within the noise by computing mean correction factors for each pixel. This includes computing the variance of the correction factors, i.e., the fitting noise. After most of the resolvable structure has been obtained, the correction-factor variance (CFV) is essentially the *square* of the population *Noise-to-Signal* ratio for the corresponding pixel, and is typically much less than one. Since each pixel has a CFV value, the CFV information may be viewed in image form. Where there is no significant signal, the CFV is generally large (approaching unity). Where appreciable signal is present, particularly noisy measurements stand out noticeably (e.g., undeglinted particle radiation). It is important to know whether such measurements affect objects of interest in the image, and the methods discussed in sections 4 and 5 below can be used to eliminate individual detector samples that may be affecting the image adversely.

The CFV image is also useful for determining the consistency in measurements obtained at different epochs. By imaging data from two epochs both separately and combined in a single image, it is possible to see where the CFV is enhanced in the combined image. A useful exercise is to compute a CFV ratio image, where the combined CFV is in the numerator and the CFV from one of the separate epochs is in the denominator. The result is related to the F-statistic, and regions where the ratio is significantly greater than unity may be where systematic incompatibilities exist between the two epochs. This may result from photometric variability or proper motion, but most often is caused by such phenomena as zodiacal background variation or hysteresis affected by changes in scan angle.

Figure 1 shows the intensity map of M31 at 100 microns after 20 iterations. The corresponding CFV map is shown in Figure 2. Both figures employ logarithmic scaling, with that of the CFV map set to show an area of enhanced fitting noise near the nucleus. The CFV is generally small in the region of the galaxy, and all but one of the peaks inside the dark area are associated with low-brightness regions. The peak just southeast of the nucleus reveals the presence of a low-level data glitch that causes a slight brightening at that location which is removable with the methods discussed in sections 3 and 4 below.

3. COVERAGE AND DETECTOR TRACK MAPS

Since the correction factors used in the MCM are mean values computed with response-function weighting, an image of the sum of the weights in each pixel shows the depth of coverage. This must be adjusted if inverse-variance noise weighting is also employed (by special request), but relative coverage density is

still obvious. The coverage maps are a useful diagnostic tool, since circumstances occasionally conspire to eliminate data unexpectedly.

Detector track maps are also now available. These are images which show the locations of the response-function centers of the detector samples used to construct the intensity map. Only the pixel containing the center is marked for each measurement. The marking is done by setting the pixel equal to the detector ID. This results in different detectors having different colors (or grey levels) when viewed with the IPAC image-processing program, SkyView (or similar program), with which one may “pick” specific pixels and thus learn the identity of the detector. Pixels not containing detector centers are zeroed out. Where two or more different detectors had measurements fall with centers in the same pixel, only the last is retained, but most images are sparse enough so that the trails of crossing detectors allow easy identification. The detector track map corresponding to Figure 1 is shown in Figure 3. The crossing of scans from the first two hours-confirmation coverages (HCONS) is obvious, and gives rise to large variations in the size and shape of the effective beam (see section 8 and Figure 8 below).

Since the MCM program allows selective exclusion of detectors by ID number, the detector track map is useful for determining such numbers. This may be for the purpose of eliminating a particularly noisy detector, for example. Another use currently being explored is to eliminate saturated point-source measurements; saturation is alleviated (there is hope that it can be essentially eliminated) without much loss of spatial information by this method. Generally some measurements see the bright object in the shoulders of the response functions, where the response functions are still fairly accurately known but the data are not saturated. If all measurements at higher response can be eliminated, the low-response measurements may construct the image correctly. When the saturation is not too severe, some relief is possible just by masking out measurements marked as associated with point sources by the data preparation program as part of its deglitching processing.

4. N-SIGMA DATA EDITING

In most iterative averaging processes, an effective way to eliminate “outliers” is to exclude samples that differ from the previous iteration’s mean by more than some threshold in units of the previous iteration’s standard deviation. This is available (applied to the correction factors) by special request in the MCM program and has been found to be quite powerful in some tests cases, but it has not yet been used extensively enough to build up a foundation of experience involving many varieties of image and settings of the n-sigma level and number of iterations to perform before beginning the editing. Furthermore, while it can eliminate practically all the flux associated with a feature caused by an outlier, it cannot generally remove all trace of the shape of the feature once it has been injected into an image, since the information contradicting the shape is usually not in the data. It may also have adverse effects on some parts of the image while fixing others. This method was used to eliminate the excess brightness of the data glitch that caused the enhanced CFV near the nucleus of M31 shown in Figure 2, as mentioned in section 2 above.

5. SCAN ANGLE AND SCAN DIRECTION EXCLUSION

Bright objects near each other in the scan direction can be affected by hysteresis, as well as other effects (e.g., “tails” in the response). The MCM program provides two ways to isolate scans. When using additional observation data, the user may specify that only scans in the survey direction be used, or alternatively, only scans in the antisurvey direction. Constructing images both ways and comparing them sheds considerable light on spurious interaction between sources. Another exclusion option is to accept only scans with angles within specified ranges. This is useful for survey data taken several months apart. These methods were used to establish that the bright object just northwest of the nucleus of M31 (see Figure 1) is real and not an artifact caused by hysteresis operating on the nucleus. Although hysteresis makes this object 7% brighter in the image in which it is downstream from the nucleus, this variation is expected and is much too small to damage the source’s credibility.

6. SIMULATION

The IPAC image construction program works by simulating the scanning of a trial image and comparing the results to the observations. A natural extension was the addition of the capability to accept an arbitrary initial image and to output the simulated measurements obtained from it. The simulated CRDD (Calibrated Reconstructed Detector Data) is realistically aliased, noise can be added to it if desired (baseline error and sample-frequency noise with an arbitrary distribution), and certain types of response-function model error can be invoked.

This capability has been employed in four activities so far: (a.) it has provided CRDD from a known scene for testing destriping algorithms; (b.) it has been used to postulate scenes that were subsequently “scanned” and images constructed, with comparison between the simulated and actual images shedding light on subtle features; (c.) it is used in creating the “beam sample” maps discussed in section 8 below; (d.) it is used in the “cross-band simulation” method of registering spatial resolution for different wavelengths.

Cross-band simulation is one of two methods for matching the resolution in one wavelength to that in another (the other, “prior knowledge”, is described in the next section and requires caution). Suppose, for example, that a color-ratio map involving 60 and 100 microns is desired. Simply computing color ratios in each pixel by employing the two intensity images typically yields fantastic artifacts because of the incompatible resolutions. The following method transforms each intensity image to obtain corresponding ones with compatible resolution without abandoning spatial information unnecessarily.

First, images are constructed in the two bands. Then the 60 micron image is “scanned” with the 100 micron detectors, using the same scan geometry and measurement locations as the real 100 micron image. The simulated CRDD is then used to construct an image. In the same way, the real 100 micron image is used to produce simulated 60 micron CRDD, which is used to construct an image. The two simulation images then have compatible resolution, and may be used to create a color-ratio map or a multi-band intensity image. Figure 4

shows a 100 micron/60 micron color ratio map for M31. In this image, brighter features have lower blackbody temperatures. Note the feature just southeast of the nucleus that is apparently very cold: this is the artifact caused by the data glitch discussed in section 2, and it does not appear when the data glitch is removed with the method of section 4.

7. PRIOR KNOWLEDGE INPUT

The IPAC MCM program actually works by modifying an initial image only at the spatial frequencies necessary to make it yield the observations to within the noise. Normally the initial image is flat, so that the final image retains as much of the initial flatness as is consistent with the data. If an arbitrary initial image is used, it too is modified only as needed to make it consistent with the observations, but the result is not generally a maximum-correlation image, because the pixels are not generally 100% correlated at the beginning. Most often it is the higher spatial frequencies that need the least modification and which therefore persist.

The use of a nonflat initial image is called "prior knowledge", since normally this is done only when one believes one has information consistent with the observational data but going beyond them. This term is used to describe the injection of external information into imaging calculations in a variety of ways. The IPAC MCM program employs prior knowledge in an extremely literal manner: it accepts a complete picture of what one believes to be the truth and proceeds to make it consistent with the IRAS data.

The prior knowledge technique has been used at IPAC as an alternative to cross-band simulation for obtaining color ratio maps. The example using 60 and 100 micron data discussed in the preceding section is handled as follows. The normal 60 micron image is used as the prior knowledge image for the 100 micron image construction, and vice versa. This amounts to assuming that if the structure in these two wavelengths could be perfectly resolved, it would be the same. Caution must be exercised because, for example, a star visible at 60 microns but not at 100 microns could lie in a region containing a cold cloud visible only at 100 microns. Using the 60 micron image as prior knowledge for the 100 micron image construction would cause the cold cloud emission to be concentrated as much as possible at the star position.

When used with care in applicable situations, this method can yield consistent resolution at two wavelengths with less overall loss of spatial information, since the resolution in the two resulting images is closer to the shorter-wavelength image, whereas cross-band simulation yields images slightly lower in resolution than the longer-wavelength image. The cross-band simulation method should always be used as a rough confirmation of the prior-knowledge method. Good results have been obtained in this way for M31; the corresponding color ratio map is very similar to Figure 4, which was produced with the cross-band simulation method discussed in the previous section.

Another use for the prior-knowledge technique is in producing super-resolved images, i.e., images with structure at spatial frequencies beyond the diffraction limit. For example, when one is certain that one is dealing with point sources (e.g., raster scans of the asteroid Egeria), the resolution is limited only by the

pixel size. This capability has been demonstrated to work only via simulation at present.

Finally, employing prior knowledge input can be used to suppress ringing around bright point sources observed against background emission. The ringing occurs because such an image satisfies the measurements with higher pixel correlation (and higher entropy) than the intuitively expected image. It is possible via an iterative procedure to estimate the point source location and amplitude and employ this information as prior-knowledge input, refining the estimates until the ringing is suppressed. The ringing pattern itself contains information, and is determined primarily by the response function size, sample density and degree of uniformity, and the ratio of the point source flux to the extended emission. Overestimates of this ratio yield a ringing pattern that is distinctively different from that of underestimates (primarily in phase), and this information can be used in the iterative procedure. This typically also yields super-resolution. The resulting lower pixel correlation (and entropy) reflects the injection of external pixel decorrelation (or negative entropy) in the form of the prior knowledge. This has also been demonstrated via simulation and is ready for scientific application. Figure 5 shows a beam-sample map (see section 8 below) of M31 at 60 microns in which the imaged spikes ring noticeably against the smoothed background (a negative print is used to make the ringing more visible). Figure 6 shows the corresponding image with ringing suppressed.

8. BEAM SAMPLE MAPS

The effective beam size in typical high-resolution IRAS images varies by factors of three over distances of several arcminutes. In order to estimate the beam size at any given position and to see typical variation over the field, “beam sample maps” are provided. These are produced from simulated CRDD based on actual coverage geometry, with the simulation scene being a collection of spike sources against a smooth background. An image of the reconstructed spikes is generated with all the same processing options as the actual image.

Three types of spikes are available: uniformly spaced spikes (the default spacing is 12 arcminutes), spikes located where the deglitching algorithm found point sources, and spikes at user-defined locations (specified either in RA and Dec or in IRAF pixel coordinates). By default only the first type is used; the pattern is phased so that there is a spike at the center of the image, but the phase can be overridden. The spikes default to fluxes of 10 Jy, but can also be specified by the user or else an auto-scaling option is available which employs the pixel flux histogram of the actual image.

By default the background is flat and equal to the actual median background. By option, the actual background can be smoothed with obvious point sources removed. It is important to have a realistic background level, since the convergence rate for point sources is affected by background. Figure 7 shows automatically scaled spikes of the first two types and smoothed background for M31 at 100 microns. The simulated CRDD obtained from this scene were used to construct the beam sample map shown in Figure 8, in which the considerable variation of effective beam size and shape can be seen, as well as the higher state of convergence of spikes in the field relative to those embedded in the extended




Fig. 5: Beam sample map in 60 microns, showing point-source ringing




Fig. 6: Same as Figure 5 except with suppression of point-source ringing




Fig. 7: Spike map with smoothed background for Figure 1




Fig. 8: Beam sample map constructed from simulation CRDD based on Figure 7

background.

An automatic analysis of the reconstructed spikes is provided in a text file. This includes the full width at half maximum on the principal axes of the beam shape, the twist angle, centroid error, and peak intensity, along with the pixel coordinates, RA and Dec, and simulated peak intensity and background.

9. IMAGING WITH A FLUX BIAS

Although the MCM program has no problem handling negative fluxes, there are reasons why positively constraining the image is desirable (e.g., reduced ringing and the inability of negative noise spikes to form), so this capability is included. On the other hand, simply excluding negative data (e.g., negative noise excursions after baseline removal) results in biased noise. The solution is to employ a flux bias, and this is now done by default.

The data preparation program informs the imaging program of the most negative intensity, and an appropriate bias is used. The bias is inserted by adding an intensity to each detector sample appropriate to the detector's response solid angle. The bias is removed from the final image by subtracting the appropriate flux from each pixel. Since the solid angle of the pixels is known exactly, but those of the detectors have estimation errors of a few percent, large flux biases cannot be used, because doing so would inject stripes into the image (in fact, this effect could be used to renormalize the detector solid angles). For most cases, this is not a problem, and the benefits of flux-bias processing far outweigh the cost. When baseline removal is used to prepare the data, a small positive flux bias is normally used. For destriping only, a negative flux bias with a larger amplitude is typical; this brings the action closer to the zero level, reducing ringing and accelerating convergence of higher spatial frequencies. Note that in the noise-to-signal interpretation of the CFV, the signal is that with the flux bias in effect.

10. SOME ADDITIONAL ENHANCEMENT TECHNIQUES

Two other techniques are worthy of mention. One of them improves the appearance of images, and one of them provides virtually instantaneous convergence at the expense of significant cosmetic degradation.

When image construction is carried to a high number of iterations, sample-frequency noise is typically driven into spikes. Since real features normally cannot be converged into spikes (with reasonable pixel sizes and typical measurement densities), and since noise features are usually accompanied by relatively high fitting variance, most of these noise spikes can be removed as follows: (a.) a spike filter (a 3×3 zero-sum filter with a central peak) is evaluated everywhere; (b.) when its output is above a threshold, the correction-factor variance of the central pixel is tested against another threshold; (c.) when both thresholds are surpassed, the pixel value is replaced with the median value of the surrounding nonspike pixels (i.e., pixels which did not also ring the spike filter). This has been found to improve image appearances and is available upon request.

Finally, an interesting technique has been found to accelerate convergence: instead of computing a mean correction factor for each pixel on the first iteration,

the minimum factor from a measurement with response above a threshold at that pixel is used. This discards flux conservation at this point, but one more iteration rescales the image, so that flux is better conserved after two iterations than normally occurs (typically about ten iterations are needed to stabilize the total flux to better than 1%). The resulting image suffers from having fit all the model errors, just as images taken to 1000 iterations have been found to do.

ACKNOWLEDGMENTS

The authors would like to thank E. L. Kopan, M. Melnyk, S. Price, and G. Veeder for useful conversations that resulted in contributions to the work described herein. This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

- Aumann, H. H., Fowler, J. W., and Melnyk, M. 1990, *AJ*, **99**, 1674
- IRAS Catalogs and Atlases: Explanatory Supplement*. 1988, ed. C.A. Beichman, G. Neugebauer, H.J. Habing, P.E. Clegg, and T.J. Chester (Washington, DC: GPO)