

Wide-field Infrared Survey Explorer (WISE)

Source Photometry Subsystem Design Peer Review Report

12-June-2008

Prepared by: R. Cutri, T. Jarrett, K. Marsh



**Infrared Processing and Analysis Center
California Institute of Technology**

WSDC D-A004

-

1 INTRODUCTION

A peer review of the design of the Source Photometry module of the WISE Science Data System was conducted by the WISE Science Data Center (WSDC) on April 4, 2008. The Source Photometry module (WPHOT) is responsible for measuring the brightness of sources detected on WISE image data. WPHOT computes instrumental brightnesses in several different ways, and these brightnesses are later calibrated by the Photometric Calibration subsystem (PCAL) for inclusion in the WISE Source Catalog. Another function of WPHOT is to compute optimized source position in pixel coordinates. These positions are converted to calibrated J2000 equatorial coordinates using the World Coordinate System (WCS) provided by the Position Reconstruction subsystem (PREX). This peer review focused on the design of the single-frame component of WPHOT.

WPHOT functionality, design and algorithms are described in Source Photometry Subsystem Design Document (WSDC D-D006).

1.1 Review Panel Members

John Krist (JPL)
Russ Laher (IPAC/SSC)
Michael Skrutskie (Univ. of Virginia, WISE Science Team)
Adam Stanford (LLNL, WISE Science Team)
Jason Surace (IPAC/SSC)
William Wheaton (IPAC/SSC)

1.2 Instructions for Review Panel

The peer review panel was asked to comment on the following specific questions:

- Does the design of the Source Photometry subsystem/module address the requirements on the system?
- Are the Source Photometry algorithms suitable and appropriate to carry out the system functions?
- Is the design robust to circumstances that will be encountered with the WISE data?

In addition, comments on other aspects of the design were welcomed.

Written reports were received from Krist, Laher, Skrutskie, Stanford, Surace and Wheaton. These reports are provided below.

1.3 Applicable Documents

WSDC Functional Requirements Document (WSDC D-R001)
WSDS Functional Design Document (WSDC D-D001)
Source Photometry Subsystem Design Document. (WSDC D-D003)

2 PANEL REPORTS

2.1 *John Krist*

Thank you for inviting me to be a WPHOT reviewer. In general, I found the plan to be well defined. Being largely based on previous 2MASS is definitely an advantage. The system as presented appears to meet spec, within the current ambiguities regarding the PSF (as described below). The algorithms appear suitable for the application, though they need to be proven on noisy model data.

Below I offer some comments and questions regarding the design spec; I refer to equation and section numbers in the document. Some of these duplicate comments that I made during the presentations.

- 2.1.1 My primary concern is the lack of knowledge of the PSF and how that may impact the assumptions on the creation of the PSF grid for profile fitting and how those profiles are applied. It was stated that the FWHM will be about 6 arcsec in the 3.3, 4.7, and 12.0 bandpasses. That's a pretty big wavelength range to have a constant-width core in a system not limited by seeing. I strongly suggest applying pressure on the project to provide reliable estimates of the PSF.
- 2.1.2 The use of pure sinc interpolation may create artifacts in the presence of airy rings, noise, and cores that are nearly-Nyquist sampled. I have found that a damped sinc interpolation provides fewer artifacts. I use a Lanczos-3 kernel sinc interpolator when resampling my PSF models.
- 2.1.3 The units of the various terms, notably flux, are not explicitly stated.
- 2.1.4 In Section 2.2, items 4 & 5 are duplicates.
- 2.1.5 It should be stated that the PSF (H) in equation 3 is normalized to a total flux of 1.0.
- 2.1.6 The equation for noise (eq. 4) has a gain term in the shot noise component but not in the flat field noise component, though both use the same flux.
- 2.1.7 Shouldn't the flux (ρ) in the noise equation be divided by the square of the gain to convert it from noise counts to noise DNs? That is, to get shot noise you take the

square root of the counts and then divide by the gain to get the noise in DN. Otherwise, you're computing shot noise on DN rather than counts.

2.1.8 In section 3.2.4, perhaps isolated bad pixels can be replaced with the corresponding ones from the model PSFs obtained from the profile fitting procedure.

2.1.9 Both "G" and "g" are used for the gain.

2.2 Russ Laher

Note – Russ kindly provided comments on the Source Detection Subsystem (MDET) Design Document (MDET) that was provided as background for this review. We include those comments here for completeness, but do not address them specifically in this report

Please find below my comments/questions from the 4/4/08 WPHOT Peer Review. In general, I am impressed with the thoroughness of the review and the sophistication of the selected algorithms. In my opinion, the design addresses the requirements on the system, the algorithms were clearly carefully selected to perform the system's functions, and robustness has been given methodical attention. I congratulate you and your team on a job well done.

2.2.1 Comments on the WPRO/WAPP charts presented at the peer review:

2.2.1.1 Maybe I missed it, but I did not see a clear formulation for how the uncertainties for the aperture photometry results will be computed (folding into it the uncertainty in the local-background estimation).

2.2.1.2 The best, unbiased estimator of the local background is the average of the pixels in the sky annulus after outlier rejection and masking off nearby sources, but if this is not possible because of imperfect outlier rejection and/or incomplete nearby-source masking, then the gaussian median is an acceptable robust compromise.

2.2.1.3 The shape of the curve of growth can determine whether nearby sources have been satisfactorily masked off. Might this be a useful diagnostic for each aperture-photometry result?

2.2.1.4 It was mentioned that the focal-plane-dependent PSFs and aperture corrections will be checked throughout the mission for stability; I recommend construction of a specific timeline for these activities during the mission (quantifying the frequency of the checking), with contingencies in case of instability.

- 2.2.1.5 After trimming lower and upper extreme values for the local-background estimation, the resulting variance will be smaller, and this needs to be taken into account. The method adopted in Spitzer downlink processing is to re-inflate the standard deviation by the deflation factor resulting from the assumption that the trimmed data are from a Gaussian distribution.
 - 2.2.1.6 Computing a source's local background for two or three different sky annuli would give users a qualitative assessment of the local-background sensitivity to sky-annulus geometry. Another interpolation method to consider for filling in missing or masked image data is the method of linear prediction coefficients (Robert Lupton of the SDSS project favors this method).
- 2.2.2 Comments on MDET SDS, v. 1.1 (1/22/08)
- 2.2.2.1 On p. 8, the noise is described as a "Gaussian" random process, yet the PowerPoint charts presented at the WPHOT peer review give a more complex noise model with a Poisson-noise component. Is this an inconsistency, or does the Central Limit Theorem apply?
 - 2.2.2.2 The detection algorithm as described in the MDET SDS is much more mathematically rigorous than that presented by Szalay et al. (1999). FYI, the following paper discusses multi-dimensional target detection more fully and includes the same formalism as the MDET SDS: Alan D. Stocker, Irving S. Reed, and Xiaoli Yu, SPIE Vol. 1305, Signal and Data Processing of Small Targets, pp. 218-231, 1990.

2.3 *Mike Skrutskie*

In general, I was left quite confident that, given the similarity to 2MASS sub-systems, the WISE WPHOT and WAPP modules have the potential to deliver source extractions that will enable catalog selection to meet the WISE survey requirements for sensitivity, completeness, and reliability. Any concerns primarily stem from fundamental differences from 2MASS. These enter in two areas, and the concerns lie downstream from the actual operation of the pipeline and more in the area of the construction of catalogs..

- 1) Active deblending - Although 2MASS pushed this element of the source extraction pipeline to the point of delivery it was not turned on and thus rigorously exercised during 2MASS. All indications are that the module is mature and will produce useful source deconvolutions. The down-the-road concern here lies the statistics of these deconvolved sources as they enter the catalog - specifically that they will be a much noisier population with potentially significant biases compared with the single source population.

Statistical uniformity will be as fundamental in the use of the catalog as temporal or spatial photometric uniformity. It will be important, either via flagging or via the construction of a separate deblend source list, to permit users to identify samples of consistent statistical heritage. A related issue - now that the system will incorporate multiple discrete PSF's across the field, it's not clear how robust the deblending procedure will be as multiple PSF's contribute to the deconvolution.

- 2) Field dependent PSF's. 2MASS used a single PSF per field, so adapting to multiple PSF's may be a challenge. For 2MASS extracting a reliable PSF that produced a consistent uncertainty was an ongoing issue and ultimately an empirical noise scaling was adopted to correct the PSF-fit value. Fortunately for WISE the PSF is better sampled than in 2MASS, so this may largely mitigate the worry that, whereas even one PSF per band was a 2MASS challenge, managing and validating multiple could be overwhelming. Of course 2MASS also had to manage seeing variability, which to first order can be considered as challenging as managing multiple spatial PSF's for a space-based system. An issue to track during initial pipeline validation will be the smoothness of transition and extracted source properties across spatial PSF boundaries. A tolerance should be established for maximal photometric offset for an infinitesimal displacement of the source across a boundary. Fortunately the PSF's (actually the focal plane positions) of the multiple apparitions of sources will be well mixed, so even a few sharp discontinuities will be mitigated by the spatial averaging. (Note that the same boundary issues will hold true for focal plane dependent curve-of-growth for the aperture photometry.)

2.4 Adam Stanford

I have read the WISE Photometry design specification document (draft version 1.1) and participated in the peer review on 4 April 2008. I believe the overall progress being made on WPHOT is satisfactory. Two important reasons for my opinion are that the design team at IPAC includes several people with a great deal of experience in the construction of the successful 2MASS photometry software, and that the design team is effectively making use of designs, algorithms, and lessons learned in the 2MASS project.

2.4.1 More specifically, I will comment as requested on three questions:

- 1) Does the design address the requirements?

Yes, the design of WPHOT clearly addresses the requirements as listed in Section 1.3 of the draft document.

2) Are the algorithms suitable and appropriate to carry out the functions?

Yes, the algorithms should be adequate to the necessary tasks.

3) Is the design robust to circumstances that will be encountered with the WISE data?

I think it is difficult to answer this question with a high degree of certainty. I would expect that because of the 2MASS background and heritage, the design would be robust. However WISE is not the same as 2MASS in several important ways, even apart from the obvious difference of being a space mission. I think it possible to depend too much on the 2MASS heritage so the WPHOT team needs to be wary when they run into problems in photometering WISE data unlike those in 2MASS.

2.4.2 Below are my other comments based on the draft document and on the peer review.

2.4.2.1 The output of the photometry routines is described as going into a text file. Is this the most efficient way to store large catalogs?

2.4.2.2 It was not clear to me how the decision will be made on the optimal size of the sky annulus. I think I agree with Tom Jarret's statement during the review that the band 4 sky annulus should be the same size as in the other bands.

2.4.2.3 I do not understand the reasoning in the way that bad pixels will be handled. It seems unreasonable to me to replace identified bad pixels, except when making "pretty pictures". Once a bad pixel has been identified then it should be flagged and not used in subsequent photometry. The datum cannot be recovered by interpolation from neighboring pixels because it never really existed.

2.4.2.4 I suspect that the PSF time variability will be more than anticipated, although it will be impossible to know the degree until in orbit data are obtained. I think the WPHOT team should try to make allowances for the possibility of having to deal with a time varying PSF.

2.4.2.5 The test plan appears to be adequate, if not daunting. I will be very interested to see the results of the tests.

2.5 *Jason Surace*

Note – Jason kindly provided comments on the WSDS Functional Design Document (FDD) and Source Detection Subsystem (MDET) Design Document (MDET) that were provided as

background for this review. We include those comments here for completeness, but do not address them specifically in this report.

2.5.1 OVERALL

First let me say I believe the overall structure of what you are doing with the WPHOT system should of course be adequate for the overall mission goals. Given IPAC's track record, I would be surprised if it were otherwise.

2.5.2 TESTING, REQUIREMENTS, and EXISTING KNOWLEDGE

However, my primary comment concerns how the requirements drive the software development. As we know, there are many ways to define flux measurements and many ways to carry out those measurements, and to a large extent they depend on the nature of the data. Now, perhaps the design studies are missing from these documents, and as a result I am getting an incomplete picture of what is happening. But as they are currently are, both WPHOT and MDET have a sense of software development for it's own sake, rather than development to solve a specific problem associated with WISE. WPHOT's simultaneous multi-band extraction represents a significant departure from the well-understood techniques used by other projects within IPAC (i.e. single-band extraction and catalog space bandmerge). I assume that at least one driver is the sheer computation time associated with catalog-based associations, which might be crippling to a pipeline trying to keep up with the satellite data flow. WPHOT should eliminate that. However, there is no cost-benefit analysis describing this.

Section 7 of the WPHOT SDS describes at least some of the necessary tests. But those analyses should be used to drive the photometer design, not just to demonstrate it's performance after the fact. For the most part, there exist Spitzer surveys of much greater depth and spatial resolution than WISE in regions that bracket pretty much the whole parameter space of the WISE survey. Key questions are:

a) Is confusion significant at these wavelengths and depths, which in turn will drive the design of the photometric extraction? This particular form of multiband extraction has been explored in detail by the GOODS team, but it was entirely driven by confusion arguments. What do high and low galactic latitudes look like? The combination of GLIMPSE, MIPS GAL, and SWIRE provide the actual source distributions in representative high and low galactic fields.

b) Are the sources likely to be point sources or not? The document mentions offhand that WISE will not resolve most objects. This is immediately addressable with the multi-degree COSMOS survey data, which also contains deep HST imaging. Recently we have come to realize that IRAC resolves pretty much all extragalactic sources brighter than 100 micro-Jy, and that point source photometry substantially under-estimates their true fluxes. This was

very much not obvious from the appearance of the images, where most things look like point sources.

c) What about the structure in the background? Background estimation is critical for extraction of faint sources. While relatively simple at high galactic latitude and at the shorter wavelengths, this becomes very complex approaching the galactic plane. Again, GLIMPSE and MIPS GAL can be used to figure this out.

In any case, I am sure you will have to look at the above anyway, since the use of deeper, higher resolution data is the only real way to properly derive completeness and reliability (C&R).

2.5.3 MDET

Optimal multiband coaddition and detection was pursued by SWIRE, namely in using the Szalay optimal multiband coadds as detection images in conjunction with SExtractor in its so-called "two-image mode". In this mode we do the detections using the multiband coadds, and that information is then used to perform source extraction from the individual bands at the locations determined by the multiband image. In the end, while the results were of some use to the science team, they had almost no impact on the survey and hence I didn't pursue the development very far (those damned requirements again). This was a result of our C&R testing and the underlying requirements of the survey release. Our C&R is derived from a 0.25 square degree portion of the survey which is roughly 10x deeper than the nominal survey depth. Our reliability requirement is very high, and in fact the desire to only release reliable sources ultimately sets our depth, which we describe based on the 95% completeness level. The use of the multiband coadds did in fact increase our completeness (as one might expect) such that in principle the survey would be substantially deeper (perhaps 20-30%). However, the reliability of those additional sources was not high enough to be useful, and if anything the use of the multiband data actually added sufficient spurious sources at relatively high SNR that its inclusion would have actually slightly raised the survey flux limits in order to maintain reliability! So we have the classic tradeoff of completeness vs. reliability - the multiband detection adds to the completeness, but you have to decide if they are good enough to use. My offhand belief is that part of the problem has to do with detector artifacts and other spurious sources spilling over from one band into another - as a result of the coaddition, every band is affected by the spurious sources in all bands.

I also note one additional thing about the use of multiband coaddition that is often lost - the actual source extraction is still done on the single band images, and as such is still ultimately limited by the SNR of that data. So although the multiband coadd may well indicate that yes, there is an object present, it does nothing to allow one to extract the flux of that object more accurately. The noise level is still what it is. For SWIRE we chose to indicate low SNR sources purely as upper limits. So one must consider if the additional sources measured as a

result of using the deep multiband coadd in fact would even have their measurements placed in the catalog.

2.5.4 PIXEL REPLACEMENT (section 3.2.4 of the WPHOT SDS)

I am not certain I agree philosophically with the idea of replacing or interpolating over missing pixels. First, frankly given the redundancy that I assume is built into the WISE survey I doubt that this circumstance will arise very frequently. But when it does, I do not agree that interpolating over pixels is the right thing to do in conjunction with aperture photometry. I tend to think of aperture photometry as the most "honest" of flux measures, since fundamentally it is just the sum of pixel surface brightnesses, which is the fundamental detector observable. Thus, it reports actual data and makes no assumptions about the underlying emission (as, for example, psf-fitting does). Thus, missing data should be treated as that, missing data. To interpolate those pixels is fine for press-release images, but not really for scientific data.

2.5.5 BACKGROUND DETERMINATION (Section 3.0.2)

There was actually not very much thought put into the fiducial aperture for IRAC. It is very large (12" radius) because we felt that was a number so large that we were pretty sure that it would capture >95% of the beam, and thus would be useful in the early part of the mission before the beam profile was adequately characterized. In fact, the ideal (minimal noise) aperture for point source aperture photometry is a diameter equal to twice the beamwidth (fwhm). This is a radius of roughly 2". This was determined (by SWIRE, anyway) empirically by examining stellar color-magnitude diagrams and choosing the apertures that minimized the observed scatter. It could similarly be derived from careful analysis of the beam profile and the detector noise properties.

Setting the background is a fairly complex issue, most approaches involve something like a modal filter, and in fact what is proposed here is essentially a median ring filter with object rejection. This should work fairly well, but I encourage an comparison versus other filters, as this can inject very complicated biases in the photometry at low flux levels. The spatial scale (i.e. width) of the filter is very sensitive to the structure of the background. Again, knowledge from Spitzer should be very useful here.

Finally, it is imperative that the method used match that used to derive the aperture corrections.

2.5.6 SymLinks (functional design section 9.2)

I am concerned about the use of symbolic links in the data storage. As is noted in the text, this is potentially a source of serious snarls in the data structure. Although in principle it allows overlaying of different hardware into the existing path structure (as might be desirable, for example, in case of hardware failure), it is also the case that symlinks fail in a

manner that is often difficult to trace. Specifically, it becomes very easy for a missing or failed or improperly made link (as might occur when swapping hardware) to result in deposition of data into places it was not intended, and this might happen differently on different computer systems which think they are linking to the same place, but in fact are not. Once such an error is made this very hard to undo, since different sets of data have now forked to different locations. At the very least, this requires very strict process control to keep track of the links.

2.5.7 Calibration

Legacy from ground-based observing is evident at numerous points in the data processing stream. This is particular notable in the apparent use of running-derivation of the flats and darks. This is very cpu-intensive, and I very much hope it is not needed. My belief is that the detector gain matrix (flatfield and flux calibration) should be extremely stable. In terms of "sky" subtraction, the most significant component is likely to be scattered earthshine (which has been a significant problem for the Akari mission).

2.5.8 DETERMINING IF AN OBJECT IS EXTENDED

Use of X^2 is one way to do this, but X^2 can be poor for very many reasons. Another way which is very effective is to look at the isophotal area. For a given dataset, a graph of isophotal area vs. aperture magnitude neatly bifurcates into extended vs. non-extended objects, as extended objects have larger isophotal areas than point sources at any given magnitude. From this one can define a parameter envelope for extended sources. We used this method in SWIRE and also analysis of iRAC Dark Field HST data. It has proven very successful. It produces nearly the same results as SExtractor's mag_best, which in turn folds in it's AI driven stellarity classifier.

Our approach in the SWIRE survey was to provide both point source and extended flux measurements, and then provided a flag to indicate which was the most appropriate "total" flux measurement.

2.5.9 GLITCHES

There is very little discussion of robustness to radiation hits and detector artifacts. They are the dominant source of unreliability in IRAC surveys such as SWIRE, and given WISE's location in LEO, I will be surprised if the radiation issue is not an order of magnitude greater. This will be doubly important when extracting photometry from the individual frames instead of the coadds - I see that WPHOT expects glitches to be masked upstream. I assume that some process during coaddition will generate masks based on some outlier algorithm and propagate these back to the individual frames that feed into WPHOT.

2.6 *Bill Wheaton*

I begin with the questions on VG46 of the presentation that you asked us to address. I hope these are the most useful things. The last few questions are just sketched in.

I also have a few general thoughts that I will add later.

2.6.1 Focal-plane dependent PSFs -- segment size? #PSFs?

See: Shupe et al 2005 <http://ssc.spitzer.caltech.edu/archanaly/shupeADASS.pdf>

I expect you have seen Dave Shupe et al.'s 2005 ADASS paper on the SIP (Simple Imaging Polynomial) convention for representing optical distortions in FITS headers. It is in the SPITZER archives, and I have attached a pdf. I do not know to what extent it has been accepted and developed for wider use beyond Spitzer. Anyhow, it simply represents the transformations from sky to detector, and the inverse, in terms of polynomials. SSC must have code in the pipeline already for populating the headers, which you can probably adapt.

I have used these for the MIPS 70 um PSF photometry and have some IDL routines that I could provide as examples if you are not familiar with this system. Unfortunately my code is specific to the 70 um Ge header and does not try to deal with a general polynomial, but the extension should be straightforward. I think you can get the polynomials from the Code V model for the optics, but I have never done that myself.

My procedures take a fits header and populate the polynomials, compute forward and reverse transformations from sky to detector pixel coordinates, and evaluate the Jacobian determinant of the distortion transformation, which is the factor by which the actual sky pixels's areas are changed due to distortion, and thus the flux conversion factors. (If this is insufficient for aperture photometry, where the circles get transformed into ellipses (or worse), you may be able to apply a linear transformation to the (x,y) coords of each pixel w/r to the star and then use the standard aperture photometry routines based on a list of pixels in the boundaries -- I am not familiar enough with those to know how and if that will work. I assume you are using something related to the old DAOPHOT, with its mysterious mmm subroutine.)

For MIPS 70um profile photometry, I simply evaluate this Jacobian at the center of the star position, and correct the fitted flux by that factor, assuming a single factor is sufficiently accurate to apply to the whole star image PSF. In my case the corrections are never more than 2% -- your mileage may vary.

The critical thing is of course the quantitative map of the distortion, both the Jacobian and the vector field, which I think you may not yet have available. I think a modest extension of the scheme I am using should be good enough to meet your 7% photometry spec, but it might be necessary to iterate once or twice to take account of second order effects if the distortion is large. This could probably be speeded up by putting things in a table, but you will need to

know the quantitative size of the effect before deciding if that is a good idea. Probably that can be deferred to later, when you have a better idea of where speed improvements are needed.

I expect that the segment size and number of PSFs will also depend on the quantitative distortion information. Probably the best you can do is to design the code to adapt to a variable segment size and number of PSFs, and take your best guess as to what will be needed when you have the Code V information. Then you will have a running system you can tune for speed or accuracy, as needed.

2.6.2 Focal-plane dependent curve of growth measurements?

See above comments re distortion. I suppose your choices are (1) to do everything on "sky pixels" (local atlas map pixels?) using the standard DAOPHOT routines, with a table indexed by X & Y deltas, centered on the star to get you from sky to detector pixels, interpolating the data values from detector to sky, but it sounds hard to me. Or, (2), my preference would be to try to modify the DAOPHOT code to work with lists of pixels (? maybe it does already?) and then populate those lists for the different apertures by mapping sky pixels back onto detector co-ordinates. Neither sounds very easy if the distortion is bad enough, but it probably isn't that awful.

2.6.3 Co-add measurements? Are they needed?

I have built my data-analysis concept structure on the belief that the operations of data accumulation and parameter estimation can be approached without bias by estimating parameters on individual images, discarding the uninteresting parameters, and averaging the physical estimates. So I am prejudiced against the use of coadds to get better statistics before extracting physics parameters, like flux estimates.

The rationale for this was clear in the context of gamma-ray analysis, where the backgrounds were large and variable, and lots of brutal selection cuts had to be imposed. It is not so clear to me that infrared analysis generally benefits in the same way, but I have gotten somewhat better repeatabilities with MIPS photometry on a DCE-by-DCE basis than obtained with other methods, so maybe my prejudice is more generally applicable. There is however one important thing that one must keep in mind, and that is that in doing weighted least-squares (or max likelihood) estimation the uncertainties on each datum must be estimated so that the expectation of the datum and the corresponding residual, (data-model), is zero.

The easiest way to do this is to compute the errors without using the data at all, so they are independent, and hence uncorrelated. Then I go on to accumulate data for quantities of interest beyond the main variable parameters essentially by accumulating the residuals of the fits. Details of this were published in "Multiparameter Linear Least-Squares Fitting to Poisson Data One Count at a Time", W.A. Wheaton, et al., Ap.J. 438:322, 1995. I doubt I can recommend this approach across the board for the WISE situation, where you want to

adapt a pre-existing system with minimum heartburn, but it may be worthwhile to bear it in mind.

2.6.4 Very bright stars & saturated stars--filter artifacts?

Basically I would do this by either ignoring the data affected (eg, cores of bright star PSFs) or de-weighting them in a way that is independent of each pixel, as outlined above.

2.6.5 Extended sources? (No plan or funds to properly address)

Here I would consider accumulating the residuals for each sky pixel (and pixel exposure time or exposure factor) into the sky atlas map. Again, to avoid bias, the weights need to be estimated from the fitted model, not from the data.

2.6.6 Confusion due to emission nebulae in Galactic plane.

Special case of #5, I guess. De-weighting the data in such regions may be one way. Another approach would be to do a spatial Fourier transform of the data, cut off the high frequencies associated with stellar PSFs, and then transform back to get an estimate of the Slowly Varying Background under the stars. This has the virtue that it can deal with modestly variable spatial structure in the BG in a natural way. One must however replace brighter stars with some kind of interpolated value.