

AWAIC: A WISE Astronomical Image Co-adder

Frank J. Masci and John W. Fowler

Infrared Processing and Analysis Center, Caltech 100-22, Pasadena, CA 91125, U.S.A.

Abstract. We describe a new image co-addition tool, AWAIC, to support the creation of a digital image atlas from the multiple frame exposures acquired with the Wide-field Infrared Survey Explorer (WISE). AWAIC includes preparatory steps such as frame background matching and outlier detection using robust frame-stack statistics. Frame co-addition is based on using the detector’s Point Response Function (PRF) as an interpolation kernel. This kernel reduces the impact of prior-masked pixels, enables the creation of an optimal matched filtered product for point source detection, and most important, it allows for resolution enhancement (HiRes) to yield a model of the sky that is consistent with the observations to within measurement error. The HiRes functionality allows for non-isoplanatic PRFs, prior noise-variance weighting, uncertainty estimation, and includes a ringing-suppression algorithm. AWAIC also supports the popular overlap-area weighted interpolation method, and is generic enough for use on any astronomical image data that supports the FITS and WCS standards.

1. Introduction

The goal of image co-addition is to optimally combine a set of (usually dithered) exposures to create an accurate representation of the sky, given that all instrumental signatures, glitches, and cosmic-rays have been properly removed. By “optimally”, we mean a method which maximizes the signal-to-noise ratio (SNR) given prior knowledge of the statistical distribution of the input measurements.

The Wide-field Infrared Survey Explorer (WISE) mission will be generating over a million exposures (or frames) over the sky. WISE is a NASA Midex mission scheduled for launch in late 2009. It will survey the entire sky at 3.3, 4.7, 12, and $23\ \mu\text{m}$ with sensitivities up to three orders of magnitude beyond those achieved with previous all-sky surveys. For details on the scientific goals, requirements, instrument and mission design (see Mainzer et al. 2005). One of the primary products from WISE is a digital image atlas that combines the multiple 8.8 s, $47' \times 47'$ frame exposures within predefined tiles over the sky. To support this, we have developed a suite of software modules collectively referred to as AWAIC for execution in the automated processing pipeline at the Infrared Processing and Analysis Center. The modules are written in ANSI-compliant C and wrapped into a Perl script, and will be made portable in the near future.

Here we review AWAIC’s co-addition algorithms, products, and extension to resolution enhancement (HiRes). It’s important to note that HiRes is not in the WISE baseline plan. It was implemented primarily to support offline research. The statistical robustness and performance of algorithms will be addressed in more detail in future papers.

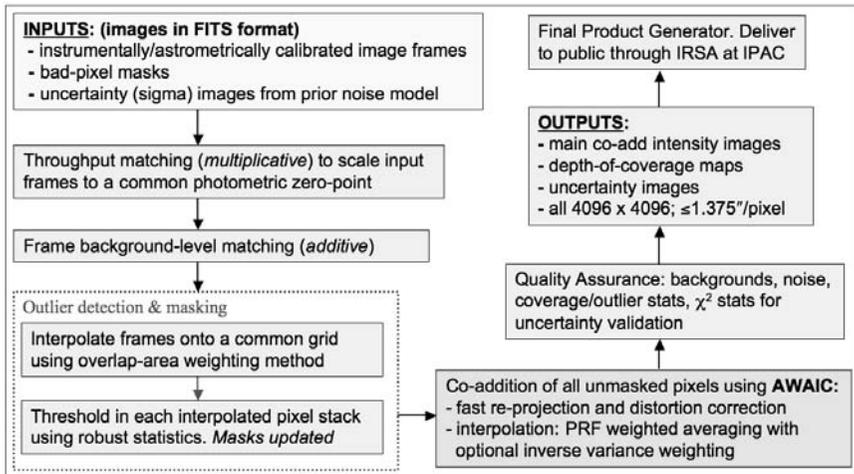


Figure 1.: Processing steps in WISE frame co-addition pipeline.

2. WISE Frame Co-addition Pipeline

Figure 1 gives an overview of the co-addition steps. The input science frames are assumed to overlap with some predefined footprint (or tile) on the sky. This also defines the dimensions of the co-add products. The uncertainty frames store $1\text{-}\sigma$ values for each pixel. These are expected to be initiated upstream from a noise model specific to the detector and then propagated and updated as the instrumental calibrations are applied. The uncertainties are used for optional inverse-variance weighting of the input measurements, and for computing co-add flux uncertainties. If bad-pixel masks are specified, a bit-string template is used to select which conditions to flag against. The corresponding pixels in the science frames are then omitted from co-addition.

The first (optional) step is to scale the frame pixel values to a common photometric zero-point using calibration zero-point information in each FITS header. Currently, the software reads a zero-point in magnitudes stored in the “MAGZP” keyword. The common (or target) zero point is then written to the FITS headers of the co-add products to enable the calibration of photometric measurements. The frame overlap (or background-level) matching attempts to equalize global levels from frame to frame, e.g., to minimize variations from instrumental transients. The method involves subtracting a “robust” planar fit to the background in each respective frame, and then adding a common level (e.g., the global median) to all the frames. This will be more or less representative of the natural background level over the co-add footprint. The goal of outlier detection is to identify frame pixel measurements of the same location on the sky which appear inconsistent with the (bulk) remainder of the sample (e.g., cosmic rays). This assumes multiple frame exposures of the same region of sky are available. Outliers are detected by thresholding the values in each aligned interpolated pixel stack using user-specified thresholds and robust estimates of the first and second moments: the median and Median Absolute

Deviation (MAD) respectively. For this method to be reliable, we require (i) typically at least five samples (overlapping pixels) in a stack, and (ii) good sampling of the instrumental PSF, i.e., at the Nyquist rate or better.

The co-addition step involves reprojecting and interpolating all “good” (unmasked) pixels onto an upsampled output grid. Details are described in § 3. The reprojection uses a fast input-to-output plane coordinate transformation that implicitly corrects for focal plane distortion if represented in the input FITS headers. The Simple Imaging Polynomial (SIP) convention for distortion is assumed (Shupe et al. 2005). The outputs from AWAIC are the main intensity image, a depth-of-coverage map, a $1\text{-}\sigma$ uncertainty image based on input *priors*, an image of the outlier locations, and optionally if the overlap-area interpolation method was used, an image of the data-derived uncertainty computed from the standard deviation in each interpolated pixel stack and appropriately scaled by the depth-of-coverage. AWAIC also produces a wealth of Quality Assurance (QA) metrics and plots over pre-specified regions of the co-add footprint. These include background noise estimates, coverage and outlier statistics, and χ^2 metrics to validate co-add flux uncertainties.

3. Co-addition using PRF Interpolation

One of the interpolation methods in AWAIC involves using the detector’s Point Response Function (PRF) as the interpolation kernel. The PRF is simply the instrumental PSF convolved with the pixel response. When knowledge of the intra-pixel responsivity is absent, the pixel response is assumed to be uniform, i.e., a top hat. The PRF is what one usually measures off an image using the profiles of point sources. Each pixel can be thought as collecting light from its vicinity with an efficiency described by the PRF.

The PRF can be used to estimate the flux at any point in space as follows. In general, the flux in an output pixel j is estimated by combining the input detector pixel measurements D_i using PRF-weighted averaging:

$$f_j = \frac{\sum_i (r_{ij}/\sigma_i^2) D_i}{\sum_i r_{ij}/\sigma_i^2}, \quad (1)$$

where r_{ij} is the value of the PRF from input pixel i at the location of output pixel j . The r_{ij} are volume normalized to unity, i.e., for each i , $\sum_j r_{ij} = 1$. This will ensure flux is conserved. The inverse-variance weights ($1/\sigma_i^2$) are optional and default to 1. The σ_i can be fed into AWAIC as $1\text{-}\sigma$ uncertainty frames, e.g., as propagated from a prior noise model. The sums in Eq. 1 are over all input pixels in all input frames. With multiple overlapping input frames, this will result in a co-add. The $1\text{-}\sigma$ uncertainty in the co-add pixel flux f_j , as derived from Eq. 1 is given by

$$\sigma_j = \left[\sum_i w_{ij}^2 \sigma_i^2 \right]^{1/2}, \quad (2)$$

where $w_{ij} = (r_{ij}/\sigma_i^2)/\sum_i r_{ij}/\sigma_i^2$. Equation 2 assumes the measurement errors (in the D_i) are uncorrelated. Note that this represents the co-add flux

uncertainty based on priors. With N_f overlapping input frames and assuming $\sigma_i = \text{constant}$ throughout, it's not difficult to show that Eq. 2 scales as: $\sigma_j \simeq \sigma_i / \sqrt{N_f P_j}$, where $P_j = 1 / \sum_i r_{ij}^2$ is a characteristic of the detector's PRF, usually referred to as the effective number of "noise pixels". This scaling also assumes that the PRF is isoplanatic (has fixed shape over the focal plane) so that $P_j = \text{constant}$. Furthermore, the depth-of-coverage at co-add pixel j is given by the sum of all overlapping PRF contributions at that location: $N_j = \sum_i r_{ij}$. This effectively indicates how many times a point on the sky was visited by the PRF of a "good" detector pixel i , i.e., not rejected due to prior-masking. If no input pixels were masked, this reduces to the number of frame overlaps, N_f .

In general, the PRF is usually non-isoplanatic, especially for large detector arrays. AWAIC allows for a list of spatially varying PRFs to be specified, where each PRF corresponds to some pre-determined region (e.g., a partition of a square grid) on the detector focal plane.

Equation 1 can be compared to the popular pixel overlap-area weighting method, e.g., as implemented in the *Montage*¹ tool. In fact if the PSF is grossly under-sampled, then the PRF is effectively a top hat spanning one detector pixel. The interpolation as described above then reduces to overlap-area weighted averaging where the interpolation weights r_{ij} become the input(i)-to-output(j) pixel overlap areas a_{ij} . Incidentally, AWAIC also implements the overlap-area weighting method, in case detector PRFs are not available.

There are three advantages to using the PRF as an interpolation kernel. First, it reduces the impact of masked (missing) input pixels if the data are well sampled, even close to Nyquist. This is because the PRF tails of neighboring "good" pixels can overlap and stretch into the bad pixel locations to effectively give a non-zero coverage and signal there in the co-add. Second, Eqs 1 and 2 can be used to define a linear matched filter optimized for point source detection. This effectively represents a cross-covariance of a point source template (the PRF) with the input data. It leads to smoothing of the high-frequency noise without affecting the point source signal sought. In other words, the SNR per pixel in the co-add is maximized for detecting point source peaks. The inclusion of inverse-variance weighting further ensures that the SNR is maximized since it implies the co-add fluxes will be maximum likelihood estimates for normally distributed data. The creation of co-adds which are also optimized for source-detection will benefit projects (e.g., WISE) where a source catalog is also a release product. The third advantage is that the PRF allows for resolution enhancement by "deconvolving" its effects from the input data.

4. Extension to Resolution Enhancement

We now describe a generic framework for co-addition with optional resolution enhancement (HiRes). Above we referred to the concept of combining frames to create a co-add. The HiRes problem asks the reverse: what model or representation of the sky propagates through the measurement process to yield the observations within measurement error? As a reminder, the measurement

¹<http://montage.ipac.caltech.edu/>

process is a filtering operation performed by the instrument’s PRF:

$$\text{sky (truth)} \otimes \underbrace{PSF \otimes \square}_{PRF} \otimes \text{sampling} \rightarrow \text{measurements.} \quad (3)$$

Our goal is to infer a plausible model of the sky or “truth”.

4.1. The Maximum Correlation Method

The HiRes algorithm in AWAIC is based on the Maximum Correlation Method (MCM). This was originally implemented to boost the scientific return of data from *IRAS* approximately 20 years ago (Aumann et al. 1990; Fowler & Aumann 1994), and is still provided as an online service to users. We have now implemented MCM in a form which is suitable for use on any imaging data that are compatible with the FITS and WCS standards, and the SIP convention for distortion. The versatility of MCM is that it implicitly generates, as its very first step, a co-add as described above. The algorithm is as follows.

1. First we begin with a flat model image of ones, i.e., a “maximally correlated” image:

$$f_j^{n=0} = 1 \forall j, \quad (4)$$

where the subscript j refers to a pixel in the upsampled output grid, and n refers to the iteration number. This starting image is a first guess at the “truth” that we plan to reconstruct. Obviously this is a bad approximation, since it represents what we know without any measurements having been used yet. We could instead have used prior information as the starting model if it was available.

2. Next, we use the detector PRF(s) to “observe” this model image, or predict the input detector measurements. Starting with $n = 1$, the predicted flux in each detector pixel i is obtained by a “convolution”:

$$F_i^n = \sum_j r_{ij} f_j^{n-1}, \quad (5)$$

where r_{ij} is the response (PRF value) of pixel i at the location of output model pixel j . Eq. 5 is a tensor inner product of the model image with the flipped PRF. It may not be a true convolution since the kernel r_{ij} may be non-isoplanatic.

3. Correction factors are computed for each detector pixel i by dividing their measured flux, D_i , by those predicted from the model (Eq. 5):

$$K_i^n = \frac{D_i}{F_i^n}. \quad (6)$$

4. For each model pixel j , all “contributing” correction factors, i.e., contributed by the overlapping PRFs of all neighboring detector pixels i are averaged using response-weighted averaging (with optional $1/\sigma_i^2$ weighting):

$$C_j^n = \frac{\sum_i (r_{ij}/\sigma_i^2) K_i^n}{\sum_i r_{ij}/\sigma_i^2}, \quad (7)$$

5. Finally, the model image pixels are multiplied by their respective averaged correction factors (Eq. 7) to obtain new refined estimates of the model fluxes:

$$f_j^n = f_j^{n-1} C_j^n. \quad (8)$$

If we are after a simple PRF-interpolated co-add, we terminate the process at step 5. In fact, Eq. 7 is analogous to the co-addition equation (Eq. 1) in that a starting model image with $f_j^0 = 1$ implies a correction factor $K_i^1 \equiv D_i$ since a PRF volume-normalized to unity predicts $F_i^1 = 1$ (Eq. 5). Therefore after the first ($n = 1$) iteration of MCM, *co-add* fluxes will automatically result: $f_j^1 = f_j^0 C_j^1 = f_j$.

If we desire resolution enhancement, the above process is iterated, where the model image from step 5 is used to re-predict the measurements in step 2. This process of iteratively refining the model continues until the model reproduces the measurements to within the noise, i.e., the predictions from Eq. 5 are consistent with the measurements D_i . If input prior uncertainties (σ_i) are available, this convergence can be formally checked using a global χ^2 test that uses all the input detector pixels:

$$\chi_n^2 = \sum_{i=1}^N \frac{(D_i - F_i^n)^2}{\sigma_i^2}, \quad (9)$$

where we expect $\chi^2 \simeq N$ (the number of degrees of freedom – the number of input pixels). Alternatively, convergence can also be checked by examining the correction factors for each detector pixel (Eq. 6), where we expect $K_i^n \simeq 1$ within the noise, or, via the averaged correction factors (Eq. 8), where $C_j^n \rightarrow 1$ after many iterations. An image of the latter can be generated by the software at each iteration. Iterating much further beyond the initial signs of convergence has the potential of introducing unnecessary (and usually unaesthetic) detail in the model. This is important to ensure a parsimonious HiRes solution.

Therefore, it is an algorithmic property of MCM that it only modifies (or de-correlates) a *flat* starting model image to the extent necessary to make it reproduce the measurements within the noise. A PRF-interpolated co-add (from the first MCM iteration) will generally not satisfy the measurements after it is “convolved” with the detector PRFs, i.e., when subject to the measurement process (Eq. 3).

In the absence of prior information for the starting model, MCM implicitly gives a solution which is the “smoothest” possible, i.e., has maximal entropy. This should be compared to maximum entropy methods (e.g., Cornwell & Evans 1985) which attempt to minimize the χ^2 of the differences between the data and the convolved model, with an additional constraint imposing smoothness of the solution. MCM requires no explicit smoothness constraint. MCM can indeed use a regularizing constraint in the form of non-flat starting model, (e.g., an image of the sky from another detector or wavelength), but this jettisons the idea of an image with maximally correlated pixels, and the refined model image will not be the smoothest possible. Smoothness is important because it can be used to convey fidelity in a model. In general, the solution to a deconvolution problem is not unique, especially in the presence of noise. Many models can be made to fit the data, and many methods invoke regularization techniques in order to select a plausible solution. A consequence is that some methods give more structure

or detail than necessary to satisfy the data, and there is no guarantee that this structure is genuine. MCM adopts the Occam's razor approach. Given no prior constraints (apart from satisfying the input data), MCM will always converge on the simplest solution. This will be the smoothest possible.

4.2. The CFV Diagnostic

A powerful diagnostic from MCM is the Correction Factor Variance (CFV). This represents the variance about the PRF-weighted average correction factor (Eq. 7) at a location in the output grid for iteration n : $V_j^n = \langle K_i^2 \rangle_j - \langle K_i \rangle_j^2$, or

$$V_j^n = \sum_i w_{ij} [K_i^n]^2 - \left[\sum_i w_{ij} K_i^n \right]^2, \quad (10)$$

where $w_{ij} = (r_{ij}/\sigma_i^2)/\sum_i r_{ij}/\sigma_i^2$, and the detector-pixel correction factors K_i^n were defined in Eq. 6. At early iterations, the CFV is generally high everywhere because spatial structure has not yet been resolved, and the model contradicts the measurements when subject to the measurement process. If after convergence, all the detector-pixel measurements contributing a non-zero response at some location j agreed exactly with their predicted fluxes (Eq. 5), then all the K_i^n would be ≈ 1 and the CFV (V_j^n) at that location would be zero. Areas with a relatively large CFV indicate the presence of input pixel measurements which do not agree with the majority of the other measurements (e.g., outliers). It could also indicate noisy data, saturated data, regions where the PRF is not a good match (e.g., erroneously broad), or that a field has not yet converged and would benefit from further iteration. By thresholding the CFV, one can therefore create a mask for a HiRes image to assist in photometry, e.g., to avoid outliers and unreliable detections from amplified noise fluctuations (see below).

Apart from providing a qualitative diagnostic, the CFV can also be used to compute *a posteriori* (data-derived) uncertainties in the pixel fluxes f_j^n in a HiRes image. In general, the $1\text{-}\sigma$ uncertainty at iteration n can be written in terms of the CFV as:

$$\sigma_j^n = c^n f_j^n \sqrt{V_j^n / \sum_i r_{ij}}, \quad (11)$$

where the sum is over the responses from all measurements i at output pixel j , i.e., the effective depth-of-coverage. c^n is a correction factor to account for re-distribution of noise power across spatial frequencies from one iteration to the next. At low iterations, power is relatively high at low frequencies, i.e., the noise is correlated across pixels. As iterations increase, noise is de-correlated and power migrates to high frequencies. The spectrum approaches that of white noise, provided the input measurement noise was spectrally white. For $n = 1$ (giving a co-add), $c^1 \equiv 1/\sqrt{P_j}$, where P_j is the effective number of noise pixels defined in § 3. With c^1 written this way, Eq. 11 becomes equivalent to the co-add pixel uncertainty defined in Eq. 2. In general, the c^n at any iteration $n \geq 1$ can be approximated from the output image products as:

$$c^n \simeq \frac{\sigma_{RMS}[f_j^n]}{\langle \sigma_j^n [c^n = 1] \rangle}, \quad (12)$$

where σ_{RMS} is the standard-deviation (or some robust equivalent) of the pixel noise fluctuations within a “source-free” stationary background region with approximately uniform depth-of-coverage in the f_j^n image. The denominator is the mean (or median) of Eq. 11 with $c^n = 1$ in the same region. At the time of writing, AWAIC only computes an image of $\sigma_j^n[c^n = 1]$, since it can be quite subjective on how the source-free stationary background is chosen. If such doesn’t exist, background fitting may be required with σ_{RMS} computed from the fit residuals. The user can then rescale the $\sigma_j^n[c^n = 1]$ image using the estimate of c^n from Eq. 12. This will give pixel uncertainties which are more or less statistically compatible with noise fluctuations in the HiRes’d image. Pixel SNRs will also be the maximum possible since MCM would have converged to the maximum likelihood estimate for data that were Gaussian distributed. With the correct value of c^n , a user then has an estimate of the flux uncertainty anywhere in the HiRes’d image, including at the location of sources. This will allow one to estimate uncertainties in source photometry. Noise correlations are also expected to be minimal in a converged HiRes image, or negligible if products were created with ringing suppression turned on (see below).

4.3. Ringing Suppression

Like most deconvolution methods, MCM can lead to ringing artifacts in the model image. This limits super-resolution, i.e., when attempting to go well beyond the diffraction limit of an imaging system. In general, ringing occurs because the reconstruction process tries to make the model image agree with the “true” scene with access to only the low spatial frequency components comprising the data. The input data are usually band-limited, and information beyond some high spatial-frequency cutoff can never be recovered. The best we can ever reconstruct is a “low-pass filtered” version of the truth, with the filter determined by the maximum spatial frequency the observations provide. This includes the finite sampling by pixels. A hard high-frequency cutoff will lead to sinc-like oscillations in real image space. The magnitude of the ringing depends on the strength of a source relative to the local background intensity level.

It is no accident that a solution with ringing is the smoothest (and simplest) solution possible with MCM. Anything smoother (with more low frequency power) will not satisfy the measurements when subject to the measurement process (Eq. 3). However, since a large number of less-smooth solutions can reproduce the observations, those without ringing are generally more desirable. Therefore, we relax our request for the smoothest image and use *prior* knowledge that the background and (desired) source fluxes are physically distinct and separable. There have been numerous approaches that have used this philosophy (e.g., Lucy 1994). In brief, the ringing suppression algorithm in AWAIC first generates an image of the slowly varying background for each input frame on some specific scale using median filtering; this is subtracted from the respective input frames to create the “source” images; negative noise fluctuations are set to zero, and a tiny positive offset added; MCM is then run on the background-subtracted images until convergence; the background images are combined and then added to the HiRes’d source-image product. This operation enforces a positivity constraint for reconstruction of the source signals. It ensures that source flux won’t ring against an essentially zero background level so power can be

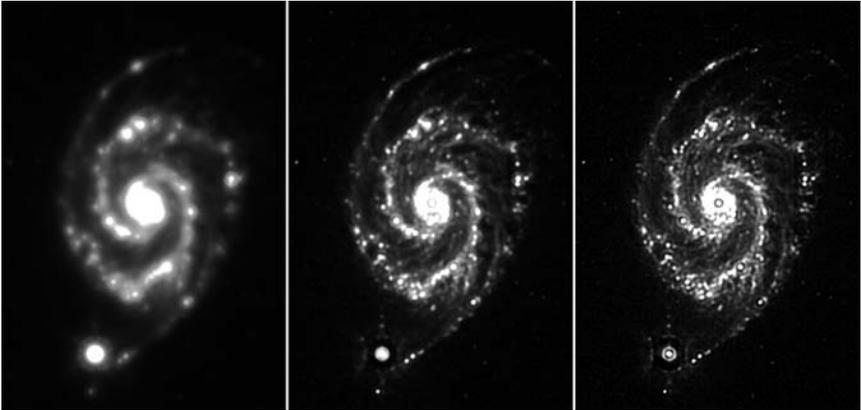


Figure 2.: M51 from *Spitzer*-MIPS $24\ \mu\text{m}$. *Left*: co-add (after one iteration); *Middle*: HiRes after ten iterations; *Right*: HiRes after 40 iterations.

forced into high spatial frequencies. The algorithm also attempts to restore the intrinsic noise properties of the HiRes process so that photometric uncertainties don't become biased. Details will be discussed in a future paper.

4.4. HiRes in Practice

Like most deconvolution methods, MCM does not alter the information content of the input image data. The signal and noise at a given frequency are scaled approximately together, keeping the SNR fixed. The process just re-emphasizes different parts of the frequency spectrum to make images more amenable to a certain kind of examination, e.g., for detecting previously unresolved objects and thereby increasing the completeness of surveys.

For optimal HiRes'ing, the input data will have to adequately sample the instrumental PSF to at least better than the Nyquist sampling frequency $2\nu_c$, where ν_c is the maximum frequency cutoff inherent in the PSF. For a simple diffraction-limited system with aperture diameter D , $\nu_c \propto D/\lambda$ and corresponds to the full width at half maximum (FWHM) of an Airy beam. Even if the detector pixels undersample the PSF (below Nyquist), redundant coverage with N randomly dithered frames can help recover the high spatial frequencies, since the average sampling will scale as $\approx 1/N$ of an input pixel. The better the sampling, the better the HiRes algorithm is at improving spatial resolution. For imaging data from the *Spitzer* IRAC and MIPS detectors with typically $\text{SNR} \gtrsim 5/\text{pixel}$ and 10 frame overlaps, our HiRes algorithm reduces the FWHM of the effective PRF to $\simeq 0.35\lambda/D$ - a factor of almost three below the diffraction limit. This corresponds to almost an order of magnitude increase in flux per solid angle for a Gaussian profile. This enhancement assumes accurate knowledge of the PRF over the focal plane.

An example output from AWaic at three MCM iteration levels is shown in Figure 2. At high iterations, point source ringing starts to appear. The ringing around the satellite dwarf galaxy at the bottom is aggravated because the core is saturated in the data, and the PRF used for HiRes'ing (which is derived from

unsaturated data) is not a good match. “Flat” core profiles in the data, due to either saturation or improperly corrected non-linearity, will contain relatively more power in the side-lobes than the actual PRF. When this PRF is used for HiRes’ing, these side-lobes will manifest as ringing artifacts in the HiRes image in order for it to reproduce the observations on “convolution” with the PRF. Even though the ringing suppression algorithm was turned on in this example, ringing is still seen around other point sources. This is because these sources are superimposed on the extended structure of the galaxy. This structure acts like an elevated background against which point sources can ring. The ringing suppression algorithm relies on accurate estimation of the local background, and this can be difficult when complex structure is involved, as it is here.

5. Summary and Future Work

We have given a broad overview of the algorithms implemented in a new generic co-addition/HiRes’ing tool. The goal is to produce high fidelity science quality products with uncertainty estimates and metrics for validation thereof. The HiRes (MCM) algorithm contains considerable improvement over previous methods in that it includes *a posteriori* uncertainty estimation, statistically motivated convergence criteria, a powerful diagnostic (the CFV) to locate inconsistencies in the input data and assess the overall quality of HiRes solutions, and the ability to handle non-isoplanatic PRFs. Algorithms will be discussed in more detail in future papers. Future work will explore methods to accelerate convergence in MCM, the ability to handle time-dependent PRFs (e.g., adapted to variable seeing), and an analysis of the completeness, reliability, and photometric accuracy of sources detected in HiRes’d images, especially in confused fields. More examples, analyses, animations of MCM, and a longer version of this paper can be found at <http://web.ipac.caltech.edu/staff/fmasci/home/wise/awaic.html>

Acknowledgments. The authors thank Roc Cutri for guidance and support. This work was carried out at the California Institute of Technology, with funding from the National Aeronautics and Space Administration, under contract to the Jet Propulsion Laboratory.

References

- Aumann, H. H., Fowler, J. W., & Melnyk, M., 1990, AJ, 99, 1674
 Cornwell, T. J., & Evans, K. F., 1985, A&A, 143, 77
 Fowler, J. W., & Aumann, H. H., 1994, in HiRes and Beyond, Science with High Spatial Resolution Far-Infrared Data, ed. S. Terebey & J. Mazzarella, 1
 Lucy, L. B., 1994, in The Restoration of HST Images and Spectra - II, ed. R. J. Hanisch & R. L. White, 79
 Mainzer, A. K. et al. 2005, Proc. SPIE, 5899, 262
 Shupe, D. L. et al. 2005, in ASP Conf. Ser. 347, ADASS XIV, ed. P. Shopbell, M. Britton, & R. Ebert (San Francisco: ASP), 491