

A Practical Demonstration of Self-Calibration of NICMOS HDF North and South Data

Richard G. Arendt¹, Dale J. Fixsen¹, and S. Harvey Moseley

*Laboratory for Astronomy and Solar Physics,
Code 685, NASA GSFC, Greenbelt, MD 20771*

`arendt@stars.gsfc.nasa.gov`, `fixsen@stars.gsfc.nasa.gov`,
`moseley@stars.gsfc.nasa.gov`

ABSTRACT

The *Hubble Space Telescope (HST)* Near Infrared Camera and Multi-Object Spectrometer (NICMOS) observations of the Hubble Deep Field-North (HDF-N) and Hubble Deep Field-South (HDF-S) are used as a test case for application of a least-squares self-calibration procedure. These data are well suited for self-calibration because the observations were dithered in a manner that provides fairly direct correlations between all the pixels of the detector. The data model we use in the self-calibration procedure explicitly includes the sky intensity, the detector gains (flat field), the detector offsets (darks), and a large-scale temporally-variable offset, in this case an offset per detector quadrant per frame which allows modeling of the “pedestal” effect of the NICMOS camera. By applying the self-calibration to the combined HDF-N and HDF-S data set, the detector properties and sky intensities are derived and their uncertainties are formally calculated. Examination of the results reveals temporal variations of the NICMOS instrument and the background sky intensity that are not directly parameterized in our data model. Comparison of the brightnesses of extracted sources with those from prior analyses shows that differences in details of the source extraction overwhelm differences between the least-squares self-calibration and other calibration of the data.

Subject headings: galaxies: general — infrared: general — instrumentation: detectors — methods: data analysis — techniques: photometric

1. Introduction

A self-calibration procedure simultaneously derives the sky intensities and detector calibration parameters, e.g. gains and offsets, from a given set of astronomical observations. There are several advantages. By using the same data to calibrate the instrument and to observe the sky, valuable

¹Raytheon ITSS

observing time can be saved. Since the same data are used for both observation and calibration, drifts and differences in the instrument between the calibration and observation are eliminated. This general approach allows calibration with a minimum of ad hoc assumptions and time consuming observations. Finally, the use of a least-squares fit in the process yields a measure of the uncertainty and of the correlations introduced by the calibration process.

Kuhn, Lin, & Loranz (1991) present a method for least-squares calibration in which the detector gains (i.e. flat field) are the only detector parameters sought. A general least-squares self-calibration procedure and a specific example in which detector gains and offsets are derived are described by Fixsen, Moseley, & Arendt (2000). Observing strategies that produce data sets well-suited to self-calibration procedures are discussed in Arendt, Fixsen, & Moseley (2000).

Fixsen et al. (2000) demonstrated use of the least-squares calibration procedure on simulated data and on *HST* NICMOS observations of the HDF-S. Since the procedure used there solved only for detector gains and offsets, other variations in the NICMOS detector response had to be accounted for in an ad hoc manner. The most obvious of these unmodeled variations is the “pedestal” effect, in which the background level (offset) of each quadrant of the detector (each read out by a different amplifier) varies from frame to frame.

The NICMOS observations of Hubble Deep Fields North and South (Thompson et al. 1999; Dickinson, 1999; Williams et al. 2000; Ferguson, Dickinson, & Williams 2000) provide a very good data set for testing our procedures. Observations over three different observing programs yield a combined data set that is suited for self-calibration despite the fact that some subsets are not when taken alone. The fact that the data come from two disjoint regions of the sky allows us to demonstrate that a data set may include observations of several targets as long as the detector and sky do not vary in modes that are not part of the model. The size of the data set allows us to challenge the computational resources available. Finally, the HDF-N and HDF-S are scientifically interesting and have been previously studied, so there are independent results to compare with our present results.

In this paper we report on results obtained using a model for the NICMOS detector that includes as model parameters gains, offsets, and additional offsets (to be varied with each frame) to account for the pedestal effect. In addition to adopting a more detailed model of the detector, the present version of the code also calculates useful statistical uncertainties on the derived detector properties and the sky intensities. With the use of real data in this demonstration, we can illustrate the strengths and limitations of the least-squares self-calibration procedure. In §2 of this paper and the Appendix, we provide details of the data that were used and the calibration performed. In §3, we discuss the results and uncertainties provided by our least-squares self-calibration procedure, while §4 compares the current results with prior work. We summarize some of the strengths and limitations revealed, and possible modifications of the code, in §5.

2. Data Processing

2.1. Data Selection

The data used in this demonstration were publicly-available archival observations of the HDF-North and HDF-South regions using NICMOS Camera 3 and its F110W filter. (Comparable data taken with the F160W filter have a slightly better signal-to-noise ratio and are not as severely undersampled with the NIC 3 camera, but as either data set would serve for testing our procedures, the choice of the F110W data was a matter of convenience.) All the data were observed in MULTI-ACCUM mode using the SPARS64 sampling. For the HDF-S, the selected data comprised the 59 images that had integration times of 1152 s (25 images) or 1472 s (34 images) and that were marked as “USED” in the NICMOS Observing Log.² The Observing Log lists another 38 “USED” images with shorter integration times, which were not used in this test. For the HDF-N, the selected data came from two different observing programs. In the first program (Dickinson 1999), the entire WFPC2 field was covered at a relatively shallow depth (9 images). The same fairly wide 3x3 dither pattern was used at each of 8 locations to cover the field. The integration time for each of the 72 images was ~ 1400 s. One image was rejected because of sources that were smeared apparently due to guiding problems. Details of these observations are found in Fruchter et al. (1998) and on the HDF-S web site.³ The second program was designed to collect deeper observations over a small subregion in the HDF-North. It used the same very tight three-step dither pattern at each of 49 pointings, set in a small 7×7 grid. Integration times for these images were ~ 900 s. Details of these observations are given by Thompson et al. (1999). Of these 147 images we selected 36 that under visual inspection appeared to have relatively low noise. In addition to the sky images, our data set included 4 of the 14 dark frames with exposure times > 1000 s, that were collected during the HDF-S observations. Dark frames are not strictly needed for processing data with the least-squares algorithm, but if they are available and truly represent the offsets and dark currents of the sky observations, then they aid the algorithm in reaching an appropriate solution. The number of images used in this demonstration was limited to keep the entire data set at a manageable size for processing with the computational resources available (§2.3).

2.2. Preprocessing

The initial processing of the MULTIACCUM data sets consisted of the following steps. Within each MULTIACCUM image, the NSAMP individual readouts and the corresponding times were read. For each pixel, an unweighted linear least-squares fit was performed to determine the slope and intercept of the sampled integration ramp. The final image was determined by the slopes for

²http://www.stsci.edu/ftp/science/hdfsouth/log_nicmos.html

³<http://www.stsci.edu/ftp/science/hdfsouth/hdfs.html>

each pixel. In the fitting, several of the first and last readouts of the detector were ignored, as listed in Table 1. In cases where the correlation coefficient was < 0.95 (or < 0.93 for the dark frames), it was assumed that a cosmic ray hit was present, and the ramp was refit for slope, intercept, and the height of a step function whose position was fixed at the location of the largest slope within the ramp. This proved to be effective at removing many of the largest cosmic rays. Cosmic rays are also identified as outliers during the least-squares calibration procedure. However, their removal at that stage necessitates complete loss of the affected data (rather than just one of several MULTIACCUM readouts), and requires iteration to prevent the initially unflagged cosmic rays from corrupting the solution.

The Observing Log for the HDF-S images provides accurate positional information for registering the images. These images are mapped into a rectangular grid in celestial coordinates. For the HDF-N, the standard header information was not sufficiently accurate for properly aligning the images. We determined the registration of these images starting with those of the wide program. At each of the 8 locations, the 9 dithered images were registered with respect to the local central image in array coordinates. This was done using a Fast Fourier Transform correlation procedure. The procedure determined x and y offsets for the strongest correlation at increments of $0.^{\circ}2$ in rotation. The optimal rotations were generally found to be close to $0.^{\circ}0$, and visual inspection indicated that the alignment of frames were accurate to the nearest pixel ($0.^{\prime\prime}2$) for nearly all sources. These images were then combined into 8 subimages. Next, each subimage was correlated with its 2, 3, or 4 neighboring subimages. These two levels of correlations defined the positions and rotations used for the wide images. The deep images were registered by correlating each of the deep program images against a preliminary mosaic made from the properly-registered wide images. [Positions of point sources in the deep portions of the HDF-N (§4.2) were later found to have random $1-\sigma$ dispersions of $\sim 0.^{\prime\prime}05$ (0.25 pixel) with respect to positions cataloged by Thompson et al. (1999).] Thus, the HDF-North images are all registered relative to a grid oriented in the array coordinates of one of the individual images {data set N4UX14HMQ at $[\alpha(J2000), \delta(J2000)] = [12^{\mathrm{h}}36^{\mathrm{m}}47.^{\mathrm{s}}66, +62^{\circ}13'02.^{\prime\prime}48]$ and roll angle(V3) = $292^{\circ}164$ }. On this grid (and in all figures), north is rotated to a position angle of 112° .

The HDF-S and HDF-N images and the dark frames are processed as a single data set. The least-squares calibration procedure assigns a sky position index, α , to each datum D^i . These sky indices are used as relative rather than absolute positions. Therefore, it is possible and convenient to process the data as if the HDF-South field were adjacent to the HDF-North field. This minimizes the blank regions in the final sky map and the numerical range of the sky indices. All pixels in the dark frames are assumed to observe the same sky pixel chosen as an otherwise unobserved corner of the resulting final image. The initial guess of the intensity of this virtual sky pixel is set to zero, and in the final result, its value should be zero within the uncertainties.

At the end of this preprocessing, each MULTIACCUM integration has been reduced to a single image that has been cleaned of the worst cosmic rays and that has been suitably registered against a defined coordinate grid. These images and their (x, y, θ) coordinates are the main input to the

least-squares calibration procedure.

2.3. Least-Squares Calibration

The least-squares calibration procedure applied to the data has been described in the papers by Fixsen et al. (2000) and Arendt et al. (2000). In those papers, the data (images) generated by a set of observations were modeled as

$$D^i = G^p S^\alpha + F^p \quad (1)$$

where i indexes the data D^i (i.e. every pixel in every image of the data set), G^p and F^p are the gain and offset for each detector pixel (indexed by p), and S^α is the sky intensity at each sky location, α . Fixsen et al. (2000) show how the least-squares solution to the set of equations defined by Eq. (1) can be calculated using the iterative evaluation of a series expansion to avoid a difficult matrix inversion.

For the present work we have modified the procedure by adopting the model of

$$D^i = G^p S^\alpha + F^p + F^q \quad (2)$$

where F^q is an additional offset term that is constant for all pixels or several fixed subsets of pixels in the detector, but can change from image to image. For the NICMOS data, the index q identifies individual quadrants and runs from 1 up to 4 times the number of images in the data set. This additional offset term allows modeling of the “pedestal” effects simultaneously with the determination of the detector gain and offset. The NICMOS pedestal effects can cause each quadrant of an image to have a different background level in different images. Because the quadrant offsets, F^q , can vary from image to image (unlike the G^p and F^p terms), they also respond to any global offset variations from image to image, such as changes in zodiacal light emission or stray light. For data from other instruments, the F^q parameter may still be useful for modeling a variable offset level. Most likely, the q index would need to be redefined to count over different, more, or fewer subsets per image than the NICMOS quadrants, but this would only affect the details of a code written to apply the procedure, not the general development presented here. In the case of the Infrared Array Camera (IRAC; Fazio et al. 1998) on the *Space Infrared Telescope Facility (SIRTF)*, F^q could be used to model the “first frame” effect (Hora et al. 2000) with the index q counting over entire images or else the four sets of alternating columns in each image that are read out by the four separate amplifiers.

With a more complex model, the least-squares algorithm is extended by dividing the solution matrix into 9 parts corresponding to G^p , F^p , F^q , and their cross terms, instead of the 4 parts required in modeling only the detector gains and offsets. This entails replication of certain parts of the algorithm with minor modifications, and in the end roughly triples the number of calculations required in each term of the series expansion. [Details of the differences between the present procedure and that described by Fixsen et al. (2000) are given in the Appendix.]

In addition to calculation of the detector and sky parameters, G^p , F^p , F^q , and S^α , the current version of our least-squares calibration algorithm can also calculate selected portions of Ψ , the covariance matrix for S^α . Fixsen et al. (2000) show how the formal calculation of the covariance matrix can be performed in the context of the least-squares calibration. However, as the entire matrix is inconveniently large, our code calculates only a selected row of the matrix, Ψ_α , which provides the covariances between the selected sky pixel, α , and all other pixels in the map. The code can also be used to calculate the diagonal of the weight matrix, which is the inverse of the covariance matrix, i.e. Ψ^{-1} . In a good data set, the off-diagonal terms of the covariance matrix will be small, and thus $\text{diag}(\Psi^{-1})$ will be $\sim 1/\sigma_\alpha^2$, where σ_α^2 is the variance for each S^α based on the formal propagation of the input uncertainties through the entire calibration procedure. Examples of these results are given in §3.1.

The initial guesses provided for the algorithm could be $G^p = 1.0$, $F^p = 0.0$, and $F^q = 0.0$, but the procedure works more quickly and effectively if better initial guesses are provided. Thus to begin with, the quadrant offsets, F^q , were initialized as the differences between the median quadrant intensities in each frame and the median quadrant intensities over all frames. These were calculated jointly for North and South fields, but separately for sky and dark frames. The initial detector offsets, F^p , were then initialized to the average value per pixel of the dark frames, after subtraction of the quadrant offsets. Next, the detector gains, G^p were initialized to the median value per pixel of the sky data after subtraction of the offset terms. Then each pixel of the gain map was divided by the median of the entire gain map. i.e. the gain map was renormalized to a median value of 1.0. When the algorithm is run iteratively to further refine a solution, the values of G^p , F^p , and F^q found by the previous iteration are used as initial guesses for the next iteration.

The initial estimate for cosmic ray flagging is that the brightest 10 sky data (or 1 dark datum) for each detector pixel are caused by cosmic ray hits. The archival data in the file h1r1523rn_msk.fits were used as a mask for bad pixels and treated as additional cosmic rays. This choice of initial flagging may blot out stars in the first pass, but helps to keep the least-squares procedure from fitting bad data. On each pass through the code, after a solution is found, the data are rechecked more carefully for cosmic ray hits, and the flagging is adjusted accordingly. Newly found cosmic rays are excluded, and in some cases, data that were previously marked as cosmic rays may be unflagged.

Initial weights were defined as

$$W^i = 1/(|D^i| + 4 |\delta D^i| + 0.1)^2 \quad (3)$$

where δD^i is the difference between D^i and a smoothed version of D^i . This definition places lower weights on the brighter sources (and their immediate vicinity), and sets an upper limit on the weights in faint regions. The factor of 4 in Eq. 3 was empirically set to reduce weights on data that are most sensitive to small positional errors and intrapixel variations. The 0.1 term sets a maximum weight that is consistent with the noise distribution of the background regions of the data. After each pass, an improved estimation of weight for each datum, D^i , was made by first

calculating the differences $\Delta^i = D^i - (G^p S^\alpha + F^p + F^q)$, and then calculating the new weights as

$$W^i = 1/[(\sigma^p)^2 + (\sigma^\alpha)^2] \quad (4)$$

where $(\sigma^p)^2$ is the variance of all Δ^i from pixel p , and $(\sigma^\alpha)^2$ is the variance of all Δ^i from sky location α . After each iteration, a cosmic ray flag was set for each D^i that led to Δ^i that was a 3σ outlier with respect to both σ^p and σ^α .

For the results presented here, we used three iterative passes through the least-square procedure, with a truncation of the series expansion at 200 terms in each pass. The calculations were performed using the Interactive Data Language (IDL) version 5.3 on a 800 MHz Pentium III running Linux PC. Processing the 42.5 MB data set required ~ 600 MB of memory and took 8.3 hr for all three iterations, in order to calculate G^p , F^p , F^q , S^α , the single row of Ψ , and the diagonal of Ψ^{-1} . The calculations are faster by a factor of ~ 3 if the model does not contain the F^q term.

3. Self-Calibration Results

The final mosaicked image of the sky produced by the standard processing of the least-squares code using the gain-offset-quadrant model (Eq. 2) is shown in Figure 1. The HDF-S region is oriented with N upward and E to the left. The larger HDF-N region has been rotated 112° counterclockwise in order to fill the image more efficiently. The stretch is linear and has been adjusted to show details in the low level background, where deficiencies in the processing are most apparent. A few scattered bright pixels are a result of incomplete removal of cosmic rays. These pixels are almost entirely confined to the edges of the mapped regions, where only a small number of repeated observations are obtained. Changes in the preprocessing of the data or in the way cosmic rays are identified and removed in the least-squares code itself could improve the cosmic ray removal.

A more obvious deficiency is the presence of bright and dark stripes that indicate one edge of the field of view of many of the individual frames. This is caused by variations in the amount of vignetting that occurs across the lower edge of the array. Because the vignetting is variable, it is not modeled by the G^p term in Eq. 2; and because it affects only an edge of the detector, it is not modeled by the F^q term in Eq. 2. Thus, as an unmodeled form of detector behavior, it leaves an impression in the resulting sky. The effect that the unmodeled vigneted edge has on the sky can be reduced by de-weighting the pixels along the edge of the detector, after the solution for G^p , F^p , and F^q has been determined, but before the final sky map is produced by applying these calibration parameters to the raw data. The resulting image is shown in Figure 2 with the same stretch as used in Figure 1.

The quadrant term, F^q , in the model is important. This is demonstrated by a comparable processing of the same data using the simpler gain-offset model described by Eq. 1. The resulting sky map (Figure 3) shows large defects, caused by the lack of modeling and removing the varying

pedestal effect, either as part of the least-squares calibration, or as part of the pre-processing of the data. The four F^q parameters for each frame tend to be correlated, and therefore individual quadrants do not appear as clearly as entire frames in Fig. 3.

Comparing results from the HDF-S data processed alone with those from the processing of the combined HDF-S and HDF-N data reveals slightly more artificial structure on large scales when the HDF-S is processed separately. The improvement in the combined processing occurs because the effective dither pattern of the combined data set samples large-scale structure better than the dither pattern of the HDF-S data alone. The least-squares algorithm has difficulties deriving the detector gains and offsets starting from initial guesses that are flat at 1.0 and 0.0, respectively, when processing either of the HDF-N data sets alone, because of the regular grid-like dither patterns which effectively sample only a small number of spatial scales.

We also explored processing that involved mapping the calibrated data onto sky pixels with smaller angular sizes than the detector pixels. Ideally, we would define the small scales right from the start (i.e. S^α in Eq. 1 or 2) and perform all calculations using these small scale pixels. However, this approach has practical difficulties in that when the data are effectively diluted across a larger number of sky pixels, there may not be enough repeated observations of all sky pixels to achieve a good least-squares solution. Recognition of cosmic rays as outlier measurements of a given sky pixel also becomes more difficult. An alternate approach is to apply the least-squares procedure as usual, using sky pixels matching the size of the detector pixels, but then map the calibrated individual frames onto a new grid of small-scale sky pixels as the last step of the process. In principle, once the calibration is determined and applied to the individual frames, any mapping/mosaicking procedure can be applied to produce the final image. The procedure we use is a simple “interlacing” method, which is equivalent to a drizzle mapping with the PIXFRAC parameter set to 0.0 (Fruchter & Hook 1998). Figures 4 & 5 show the HDF-S and Thompson et al. portion of the HDF-N mapped onto sky pixels that are 0.36974 and 0.5 times the size of the $\sim 0''.2$ NIC3 pixels. These finer scale maps were produced to facilitate certain quantitative comparisons with previous results (§3.2). Figure 4b shows a detailed look at the large bright HDF-S galaxy (HDFS J223249.56-603835.5) with the default $0''.2$ pixel scale (from Fig. 2), as remapped on smaller scale pixels (from Fig. 4a), and as mapped using the drizzle algorithm (STScI). The spatial structure of the fine-scale mapping and the drizzle result are very similar, although our mapping procedure does leave some sky pixels unobserved, especially in the low coverage regions near the edge of the map. The non-zero PIXFRAC parameter used in the drizzle result fills these holes by effectively applying a slight smoothing or convolution to the data.

3.1. Formal Weights

The uncertainties are also calculated as part of the least-squares fit model. In principle this results in a $\Gamma \times \Gamma$ covariance matrix (Ψ), where Γ is the number of sky pixels observed. In the present example $\Gamma = 734758$. This matrix is too large to display (or calculate) in its entirety, but

we can easily calculate a single, typical row of the matrix and the diagonal of the inverse or weight matrix. The diagonal of the weight matrix (Ψ^{-1}) is a key to the overall uncertainties in the result. Figure 6 shows these weights using a square root scaling to provide a more useful dynamic range. It is expected that the weights, $\text{diag}(\Psi^{-1})$, should be proportional to $M^\alpha G^p / S^\alpha$ where M^α is the total integration time on a sky pixel α , G^p is the detector gain at pixel p , and S^α is the sky intensity. The effects of each of these three components are visible in Figure 6. The total coverage M^α is generally the dominant factor in the weights, creating higher weights for the deep observations of the HDF-S and the Thompson et al. portion of the HDF-N. The influence of the gain G^p appears as higher weight regions of the upper right of the HDF-S and the Thompson et al. HDF-N regions, and the lower left of the other HDF-N regions. The effect of the sky intensity S^α is apparent in that the locations of bright sources have relatively low weights, owing to the larger absolute (not relative) uncertainties on the bright sources. Cosmic rays and bad pixels also introduce pixels of low weight. These are more difficult to discern, however, because they usually affect only a single integration.

The full parameter weight matrix can be written as $\Psi^{-1} = HWH$, where W is the data weight matrix, and H is the model, e.g. $G^P S^\alpha + F^p$ (Appendix; Fixsen et al. 2000). The parameter weight matrix is dimensionally large but very sparse, whereas the covariance matrix is also large and not at all sparse. Any analysis requiring use of the parameter weight matrix can be accomplished using the component pieces W and H , each being a sparse matrix with a total of non-zero elements comparable in size to the data set being processed.

The off-diagonal components of the weight matrix or the covariance matrix can be important in some cases. By combining the images of HDF-N and HDF-S we have a sufficient amount of well-dithered data to make a good determination of the calibration. Hence the off-diagonal components are not important for this analysis. This is shown in Figure 7, which shows a complete row of the covariance matrix. As each column corresponds to a complete set of sky pixels $\alpha = 1, \dots, \Gamma$, a full row of the matrix can be formatted as an image of the sky. The dither pattern can be observed in the off-diagonal elements of the covariance matrix. Bright features also appear, as these provide most of the weight of the data to determine the gain. While the off-diagonal elements are not significant for our analysis here, if the data were analyzed for correlations, then the dither patterns will appear in the result at some level. The image here can provide an estimate of the size of these biases and can be used to quantify the uncertainties.

4. Comparison with Prior HDF Results

The HDF-N and HDF-S data sets are well-suited for testing the least-squares processing because they contain sufficiently dithered images and because the published results are available in quantitative forms that allow direct comparison with results derived from the least-squares calibration. The HDF-S data are available as raw and calibrated individual images from the STScI HDF-S

website.⁴ The final drizzled images of the HDF-S field are also available. The HDF-S source catalog generated from SExtractor (Bertin & Arnouts 1996) and the relevant SExtractor parameters are obtained from the web site. For the HDF-N, the source catalogs generated by Thompson et al. (1998) have been obtained from the Astronomical Data Center (ADC).⁵

Based on the format and availability of data for comparison, we have examined differences between our least-squares calibration and other results at several different stages. First, pixel-by-pixel comparison of the final mosaicked images provides information on the differences between the processing in the images that would typically be used for scientific analysis. Comparison of calibrated individual frames is a more direct indicator of difference in the calibration, since in the mosaicked images the effect of some calibration defects may be averaged away. Finally, a comparison of the fluxes of sources extracted from the mosaicked images shows the differences in the final product of the data analysis (for many scientific projects), although by this level of processing, the calibration of the data may not be the only aspect of the data reduction leading to differences in the results.

4.1. HDF-S

4.1.1. *Images*

For comparison with the HDF-S images in Figures 1 & 2, the STScI-calibrated individual frames were mosaicked into a sky image with identical 0''.2 pixels using the same mapping procedure that was applied to the least-squares calibrated data. Correlation plots of the images produced from the STScI data and those shown in Figure 1 are shown in Figure 9. The correlation is depicted on a logarithmic scale in order to clearly show the comparisons at all intensities. The figure also shows the correlation using only the data that are assigned high weights by the STScI processing. In these plots, a tight linear correlation with a slope of 1.00 indicates good agreement of the results. The intercept of the correlation is not important because the least-squares procedure provides only a relative calibration (i.e. the intercept depends on the absolute calibration of the data sets, which requires the observation of standard reference stars). Comparison of the two plots in Figure 9 shows that most of the residual cosmic rays that appear as bright pixels in Figure 1 are assigned low weights. In a similar comparison, reweighting of the vignetted edge before the final mapping (i.e. Fig. 2) has only small effects on pixels that are near the noise level.

The top half of Figure 10 shows the mosaicked image (on 0''.2 pixels) of the STScI calibrated data and the HDF-S portion of Figure 2. The differences between the STScI calibrated data and the HDF-S images from Figs 1 & 2 are shown in the lower half. The STScI calibrated data and our

⁴<http://www.stsci.edu/ftp/science/hdfsouth/hdfs.html>

⁵<http://adc.gsfc.nasa.gov>

least-squares calibrated data are mapped using the same procedure and image registrations for this comparison. Thus, artifacts in the difference images are not caused by any positional or geometric effects. The dark bars in upper regions of the lower left difference image are due to differences in the weighting of the data sets. Applying a deweighting to the vignetted edge of the detector makes the two sets of weights much more similar as shown in the lower right difference image. The vertical stripes running through the bright star in the upper left, and some fainter stripes at other locations, are due to electronic ghosting, which was corrected in the STScI processing by subtracting median values from affected columns, but is not appropriately modeled by Eq. 2 in the least-squares self-calibration procedure. Differences in the cosmic ray rejection techniques causes some graininess in the low-coverage regions around the edges of the difference images. Our cosmic-ray rejection also tends to reject more points in the bright sources in the data, leading to the white sources on the left side of the difference images and the systematically lower intensities for several bright pixels seen in Fig. 9. Remaining source and background variations (including the faint diagonal stripes in the upper left of the difference images) are at levels of a few percent as seen in Fig. 9, and are due to the differences between the gains and offsets derived with the least-squares self-calibration and those applied in the STScI processing.

Examination of the calibrated individual frames shows that the least-squares processing gives better results for some frames than others. The qualitative indicator of a poor calibration is the clear residual of the gain map visible in the background of the calibrated image. Figure 11 shows the raw data, the least-squares calibrated frame, the STScI calibrated frame, and the difference between the latter two, for two frames in which the least-squares calibration appears to work well and poorly. Data that are flagged in the least-squares processing as cosmic ray effects or in the STScI processing as bad data are masked in white. The main difference between the least-squares and the STScI calibrations for the good frame is the striping along the columns. For the frame poorly calibrated by the least-squares procedure, additional artifacts associated with the gain map and the quadrant offsets are also present. Figure 12 shows the correlations (as in Fig. 9) of the good and bad single frames and their corresponding STScI calibrated frames. In this representation, the differences between the good and bad calibrations is mostly evident as increased variance at low intensities.

The fact that some frames are better-calibrated by the least-squares procedure than others indicates that the model of the data (Eq. 2) is insufficient for describing the full variety of the NICMOS detector behavior. Equation 2 contains the apparently erroneous assumption that the same gain and offset images can be appropriately applied to all the images of the data set. The variations of the vignetted edge of the detector is clear evidence of this oversimplification. The more subtle differences revealed in the poorly calibrated frames seem to indicate that the apparent sky brightness is not constant as the model assumes. This is based on the observation that the frames that appear to be badly calibrated by the least-squares procedure are the ones that have systematically large values for their quadrant offsets compared to others in the data set. A simple and strong diagnostic of the badly calibrated frames is that the sum of their quadrant offsets is

> 0.037 ADU. The large quadrant offsets are derived because they are the best parameters the model can adjust to a sky which has changed in brightness, but since the quadrant offsets are not modified by the detector gain as the sky is, they cannot perfectly compensate for a changing sky brightness. In turn, this changing sky brightness appears to be caused by scattered earthlight. Figure 13 shows the correlations of the sum of the quadrant offsets with the SUN_ALT keyword from the FITS headers of the data sets. For both the Thompson et al. HDF-N observations and the HDF-S observations, the badly calibrated frames are associated with the largest elevation of the Sun above the earth limb. There is no such correlation with the values of the SUNANGLE keyword, indicating that the stray light is not directly related to the Sun itself, though it does need to be above the horizon in order for earthlight to be present. A formal solution for this problem would be to extend the data model to a form such as

$$D^i = G^p(S^\alpha + F^x) + F^p + F^q \quad (5)$$

where the new offset term, F^x , is a function of frame number, time, or the SUN_ALT value. In practice, this model may have difficulties with degeneracy between the F^p , F^q , and F^x terms.

4.1.2. Source Catalogs

We have run the SExtractor source extraction program (version 2.2.1; Bertin & Arnouts 1996) on three images of the HDF-S that we produced: the least-squares calibrated frames mapped onto $0''.2$ and $0''.075$ pixels, and the STScI calibrated frames mapped onto $0''.2$ pixels. As a control test, we also ran SExtractor on the drizzled STScI image with $0''.075$ pixels. In all cases the SExtractor parameters were set to reproduce as closely as possible those used for the STScI catalog. The extracted source lists from each of these four images were compared against the STScI list. Sources from the different catalogs that were positionally coincident within $0''.1125$ ($= 1.5 \times 0''.075$ pixels) were considered to be the same sources. The magnitude differences between sources in our catalogs and those in the STScI catalog, as measured in a $0''.5$ aperture, are plotted in Figure 14. The magnitude differences in the control test show a $1-\sigma$ dispersion of ~ 0.07 mag between our application of SExtractor and the STScI application of SExtractor on the same data. We attribute this level of variation to different versions of the SExtractor software and to slight differences in the configurations in which the software was run. In the comparison between magnitudes derived from the STScI data mapped on $0''.2$ pixels and the STScI catalog magnitudes, we find that the low resolution pixelization introduces another $1-\sigma$ dispersion of ~ 0.08 mag, for a total of $(0.07^2 + 0.08^2)^{0.5} = 0.11$ mag. Comparison between the $0''.075$ pixel data sets indicates that the differences in the data and calibration introduce a $1-\sigma$ dispersion of ~ 0.18 mag, which combines with the SExtractor variations for a total dispersion of $(0.07^2 + 0.18^2)^{0.5} = 0.19$ mag. This ~ 0.18 mag variation can be accounted for by the fact that the subset of the HDF-S data that we use includes only $\sim 73\%$ of the total integration time used by the STScI analysis. Thus, the photon counting statistics are expected to introduce a variation of $-2.5 \log(0.73^{0.5}) = 0.17$ mag, and any variation arising from differences in the calibration procedures must be small in comparison. The various

uncertainties and their sources are only approximate, as they are obtained by direct measurement from a relatively small sample. Magnitudes measured using larger or smaller apertures or by isophotal analyses show larger differences than shown here. In our present analysis, while the magnitudes of extracted sources are of scientific interest, the magnitudes of the extracted sources depend more on the details of the source analysis than those of the calibration. Therefore, we believe that direct comparison of pixel intensities (§4.1.1) is a better indicator of calibration differences than the comparison of extracted source fluxes.

4.2. HDF-N

4.2.1. *Source Catalogs*

For the HDF-N we performed comparisons of extracted source photometry in a similar fashion to those performed for the HDF-S. Source magnitudes extracted from least-squares calibrated maps at $0.^{\prime\prime}2$ and $0.^{\prime\prime}1$ pixel scales (Figs. 2 and 5, respectively) were compared to the magnitudes listed in the catalog of Thompson et al. (1999, their Table 4). The differences in magnitudes are shown in Fig. 15. The variances here are larger than they were for the comparable HDF-S comparisons by an effective additional noise term with a $1-\sigma$ amplitude of ~ 0.22 mag. Part of this additional noise should arise from fact that, in our maps, the number of images in the deep portion of the HDF-N is only $\sim 75\%$ that of the HDF-S, and the total integration time is only $\sim 60\%$ of that in the HDF-S region. Additional variations may arise from the use of SExtractor on our least-squares calibrated images versus the (predominantly) KFOCAS derived magnitudes in the Thompson et al. catalog, as is suggested by the comparison of our Fig. 15 with Fig. 8 of Thompson et al. (1999).

4.3. Source Counts

A plot of the differential source counts as a function of AB magnitude is shown in Figure 16, for sources extracted from our combined HDF-N and HDF-S data set. For comparison, results from the STScI analysis of the HDF-S, and the Thompson et al. analysis of the HDF-N are also plotted. We have not attempted to differentiate between stellar and galactic sources in any of these counts. In all cases isophotal AB magnitudes are used. For our results and the the STScI results the “corrected” isophotal magnitude provided by SExtractor (Bertin & Arnouts, 1996) are used. Our source counts do not include any sources from regions where there were fewer than 5 overlapping exposures, in order to exclude some probable false detections in these noisier regions. We have used the derived weight map (Fig. 6) to evaluate the effective area of the sky in which sources of a specified magnitude may be detected. Detection of sources in SExtractor is based on a surface brightness threshold, so that the effective survey area is a function of the size of the source for a fixed magnitude. Fig. 16a shows the survey area for assumptions of 3 different source areas, corresponding to 3.26, 6.03, and 18.45 pix^2 . These areas are the 10th, 50th, and

90th percentiles of the distribution of πab for sources with $m_{AB} \leq 24.5$, where a and b are the semi-major and semi-minor axes reported by SExtractor. These curves are directly proportional to a cumulative histogram of the noise, $\sim \text{diag}(\Psi^{0.5})$, but displaced in magnitude for different source sizes. Figure 16b shows the source counts as normalized for the same source sizes. For sources brighter than ~ 25 th magnitude, the smaller error bars on our results reflect the decrease in the Poisson uncertainties, which is obtained by the inclusion of the wider, shallower portion of the HDF-N. At the fainter end, the error bars are extended to cover the entire range implied by the assumptions of the 10th–90th percentile of source sizes. Our results are restricted to a brighter limiting magnitude for the following reasons: the counts shown in Fig. 16b are extracted from the 0.2 pixel images, in which the large pixel size tends to wash out the faintest sources; we use only a fraction of the data that are available in the HDF-S and the Thompson et al. (1999) HDF-N observations; and the present least-squares calibration procedure does not model transient changes in the detector behavior other than variable offsets in each quadrant, which leaves low level artifacts in the images that interfere with the detection and photometry of faint sources (§4.1.1).

5. Summary

The successes of the least-squares self-calibration procedure that have been demonstrated are: 1) The procedure calibrates observational data using little or no separate calibration data. 2) The procedure can calibrate detector parameters in addition to fixed gains and offsets, as demonstrated by the inclusion of the time-variable offset (“pedestal effect”) modeled here. 3) Formal uncertainties of the results are generated, for use with subsequent scientific analyses. 4) Processing reveals anomalies and defects when there are changes in the instrument or sky in modes that cannot be fit with the adopted model. 5) Data sets can be composed of observations of disjoint patches of sky, or of data widely separated in time, provided there are no systematic variations of the sky or instrument that are not modeled. 6) Combined processing of well-dithered data with poorly-dithered data can be used to establish a self-calibration across a less than ideal data set, as obtained by combining the wide HDF-N observations with the deeper HDF-S observations. 7) The procedure can produce sub-pixel resolution sky maps in cases where the coverage is sufficient (e.g. Figs. 4 & 5).

Limitations that are fundamental to the procedure are: 1) The calibration will account only for instrumental effects that are included in the model. 2) The convergence of the solution may be problematic. This is especially true when the model parameters are nearly degenerate, as is the case for our present model with two kinds of detector offsets. 3) True sub-pixels resolution, where the model for a single datum depends on the sky intensity at several locations, i.e. $H^i = G^p \sum_{\alpha} S^{\alpha} + F^p + \dots$, is major complication. Inclusion of latent images in the data would require a similarly unwieldy model. However, we note that these cited difficulties are problems that will affect any calibration process and are not unique to self-calibration procedures.

Another set of limitations is a consequence of the specific implementation of our present code.

Future development efforts will be aimed at the following goals: 1) Generalizing the structure of the code, which is currently tailored to the data model in use. Because of this, changing models is not as simple as replacing a single function or subroutine in the code. 2) Eliminating the requirement of accurate positional information for the observations. In principle, the pointing information of the instrument could be part of the data model and obtained as part of the least-squares solution, however this would be a major complication to the procedure. 3) Improving the recognition and rejection of cosmic rays and other bad data. 4) Improving the efficiency with respect to both time and memory usage.

We thank H. Teplitz, J. P. Gardner, and H. Bushouse for assistance on various elements of this work. We thank E. Fixsen for efforts aimed at improving the grammar and clarity of this paper. We acknowledge funding from NASA Office of Space Science under the *SIRTF* IRAC project.

A. A Least-Squares Calibration for Gains, Offsets, and Quadrants

In this appendix, we point out the differences between the least-squares calibrations using the detector model presented by Fixsen et al. (2000) and the more complex detector model here. This serves as example of how the procedure can be modified for various models. The notation used here is the same as that used by Fixsen et al. (2000).

In Fixsen et al. (2000), the data, D^i , are compared to a model, H^i , which is defined as

$$H^i = G^p S^\alpha + F^p \quad (A1)$$

where G^p and F^p are the detector gains and offsets for pixel p , and S^α is the sky intensity at position α .

The model parameters are found by linearizing the model, $H^i = H^i(0) + H_\mu^i \delta^\mu$, and then applying a least-squares solution to calculate the perturbations

$$\delta^\mu = \{\delta G^1, \dots, \delta G^p, \delta F^1, \dots, \delta F^p, \delta S^1, \dots, \delta S^\alpha\} \quad (A2)$$

of the model parameters from those of the initial guess $H^i(0)$. With the introduction of the weight matrix, W_{ij} , and the definition of $\Delta^i = D^i - H^i(0)$, the solution for δ^μ is given by

$$\delta^\mu = (H_\mu^i W_{ij} H_\nu^j)^{-1} H_\nu^k W_{kl} \Delta^l. \quad (A3)$$

To deal with the matrix inversion required for the solution, Fixsen et al. (2000) divided the parameter weight matrix $H_\mu^i W_{ij} H_\nu^j$ into four parts (or three independent parts)

$$H_\mu^i W_{ij} H_\nu^j = \begin{pmatrix} A & B \\ B^T & C \end{pmatrix} \quad (A4)$$

where A relates to the detector parameters, C relates to the sky parameters, and B contains the connections between the two set of parameters. Then, the inverse of this matrix can be split into similar parts:

$$(H_\mu^i W_{ij} H_\nu^j)^{-1} = \begin{pmatrix} Q & R \\ R^T & \Psi \end{pmatrix} \quad (\text{A5})$$

where

$$Q = (A - BC^{-1}B^T)^{-1} \quad (\text{A6})$$

and

$$R = -QBC^{-1}. \quad (\text{A7})$$

Then with the definition of

$$Y = H_p^i W_{ij} \Delta^j - BC^{-1} H_\alpha^i W_{ij} \Delta^j, \quad (\text{A8})$$

the solution for the detector parameters $X^p = \{\delta G^1, \dots, \delta G^p, \delta F^1, \dots, \delta F^p\}$ is obtained as

$$X = QY \quad (\text{A9})$$

$$= (A - BC^{-1}B^T)^{-1}Y \quad (\text{A10})$$

$$= A^{-T/2} (I - A^{-1/2} BC^{-1} B^T A^{-T/2})^{-1} A^{-1/2} Y \quad (\text{A11})$$

$$= A^{-T/2} (I - TT^T) Z \quad (\text{A12})$$

$$= A^{-T/2} \sum_{n=0}^{\infty} (TT^T)^n Z \quad (\text{A13})$$

where $T = A^{-1/2} BC^{-1/2}$. In practice the final form of this solution is evaluated iteratively for Z

$$Z_{n+1} = Z_0 + TT^T Z_n \quad (\text{A14})$$

and then $X = A^{-T/2} Z$. Implementation of this solution is practical because the key pieces A , B , and C are easily calculated, and because C is diagonal and A consists of 4 diagonal pieces, so their inverses are fairly simple.

Now in the present work, for processing the *HST* NICMOS data we have expanded the model to include an offset that is constant in each quadrant of the detector, but differs from frame to frame, F^q , where the index q counts over each of the quadrants in each of the frames of the data set. Thus our model is written as

$$H^i = G^p S^\alpha + F^p + F^q. \quad (\text{A15})$$

The solution for this model can be written in a manner similar to the previous one. Here the parameter weight matrix can be divided into the following parts

$$H_\mu^i W_{ij} H_\nu^j = \begin{pmatrix} A & K & B \\ K^T & J & L \\ B^T & L^T & C \end{pmatrix} \quad (\text{A16})$$

(cf. Eq. A4) where A , B , and C are the same as before, and the new parts J , K , and L contain the relations among the F^q , and between F^q and $\{G^p, F^p\}$ and S^α respectively. J is diagonal. K and L resemble B , but are more compact.

Thus, the solution for this model can be expressed analogously to that of the former model with the substitutions of

$$A \rightarrow \begin{pmatrix} A & K \\ K^T & J \end{pmatrix} \quad (A17)$$

and

$$B \rightarrow \begin{pmatrix} B \\ L \end{pmatrix}. \quad (A18)$$

With these substitutions, Eq. A9 - A13 can be rewritten as

$$X' = \left[\begin{pmatrix} A & K \\ K^T & J \end{pmatrix} - \begin{pmatrix} B \\ L \end{pmatrix} C^{-1}(B^T L^T) \right]^{-1} Y' \quad (A19)$$

$$= \begin{pmatrix} A^{-T/2} & 0 \\ 0 & J^{-T/2} \end{pmatrix} \left[I + \begin{pmatrix} A^{-1/2} & 0 \\ 0 & J^{-1/2} \end{pmatrix} \begin{pmatrix} 0 & K \\ K^T & 0 \end{pmatrix} \begin{pmatrix} A^{-T/2} & 0 \\ 0 & J^{-T/2} \end{pmatrix} \right]^{-1} \begin{pmatrix} A^{-1/2} & 0 \\ 0 & J^{-1/2} \end{pmatrix} Y' \quad (A20)$$

$$= \begin{pmatrix} A^{-T/2} & 0 \\ 0 & J^{-T/2} \end{pmatrix} \left[I + \begin{pmatrix} 0 & U \\ U^T & 0 \end{pmatrix} - \begin{pmatrix} T \\ \tau \end{pmatrix} (T^T \tau^T) \right]^{-1} \begin{pmatrix} Z \\ \zeta \end{pmatrix} \quad (A21)$$

$$= \begin{pmatrix} A^{-T/2} & 0 \\ 0 & J^{-T/2} \end{pmatrix} \left[I - \begin{pmatrix} TT^T & T\tau^T - U \\ \tau T^T - U^T & \tau\tau^T \end{pmatrix} \right]^{-1} \begin{pmatrix} Z \\ \zeta \end{pmatrix} \quad (A22)$$

$$= \begin{pmatrix} A^{-T/2} & 0 \\ 0 & J^{-T/2} \end{pmatrix} \sum_{n=0}^{\infty} \begin{pmatrix} TT^T & T\tau^T - U \\ \tau T^T - U^T & \tau\tau^T \end{pmatrix}^n \begin{pmatrix} Z \\ \zeta \end{pmatrix} \quad (A23)$$

where

$$X' = \{G^1, \dots, G^p, F^1, \dots, F^p, F^1, \dots, F^q\}, \quad (A24)$$

$$Y' = \begin{pmatrix} H_p^i W_{ij} \Delta^j - BC^{-1} H_\alpha^i W_{ij} \Delta^j \\ H_q^i W_{ij} \Delta^j - LC^{-1} H_\alpha^i W_{ij} \Delta^j \end{pmatrix} \quad (A25)$$

and

$$Z' = \begin{pmatrix} A^{-1/2} & 0 \\ 0 & J^{-1/2} \end{pmatrix} Y'. \quad (A26)$$

Analogous to T , the matrices τ and U are defined as $\tau = J^{-1/2} LC^{-1/2}$ and $U = A^{-1/2} K J^{-1/2}$.

Finally, Eq. A14 can be rewritten as

$$Z'_{n+1} = Z'_0 + \begin{pmatrix} TT^T & T\tau^T - U \\ \tau T^T - U^T & \tau\tau^T \end{pmatrix} Z'_n. \quad (A27)$$

This expression can be computed much like the former expression for Z_n , though in this case 6 matrix \times vector multiplications are required per iteration for Z'_n instead of only 2 matrix \times vector multiplications required for each iteration of Z_n . Consequently, deriving the solution for G^p , F^p , and F^q takes about three times as long as solving for only G^p and F^p .

REFERENCES

Arendt, R. G., Fixsen, D. J., & Moseley, S. H. 2000, *ApJ*, 536, 500
Bertin, E. & Arnouts, S. 1996, *A&AS*, 117, 393
Dickinson, M. 1999, After the Dark Ages: When Galaxies Were Young (the Universe at $2 < z < 5$),
ed. S. S. Holt & E. P. Smith, (Woodbury, NY: AIP), 122
Fazio, G. G., et al. 1998, *Proc. SPIE*, 3354, 1024
Ferguson, H. C., Dickinson, M., Williams, R., 2000, *ARA&A*, 38, 667
Fixsen, D. J., Moseley, S. H., & Arendt, R. G. 2000, *ApJS*, 128, 651
Fruchter, A. S. & Hook, R. N. 1998, preprint, (astro-ph/9808087)
Kuhn, J. R., Lin, H. & Loranz, D. 1991, *PASP*, 103, 1097
Thompson, R. I., Storrie-Lombardi, L. J., Weymann, R. J., Rieke, M. J., Schneider, G. Stobie, E.,
& Little, D. 1999, *AJ*, 117, 17
Williams, R. E., et al. 2000, *AJ*, 120, 2735

Table 1. NICMOS MULTIACCUM Data Fitting

Data Set	Number of Ramp Samples Rejected		Corr. Coeff. for Refitting with Step Function
	Beginning	End	
South	2	0	< 0.95
North	3	1	< 0.95
Dark	3	1	< 0.93

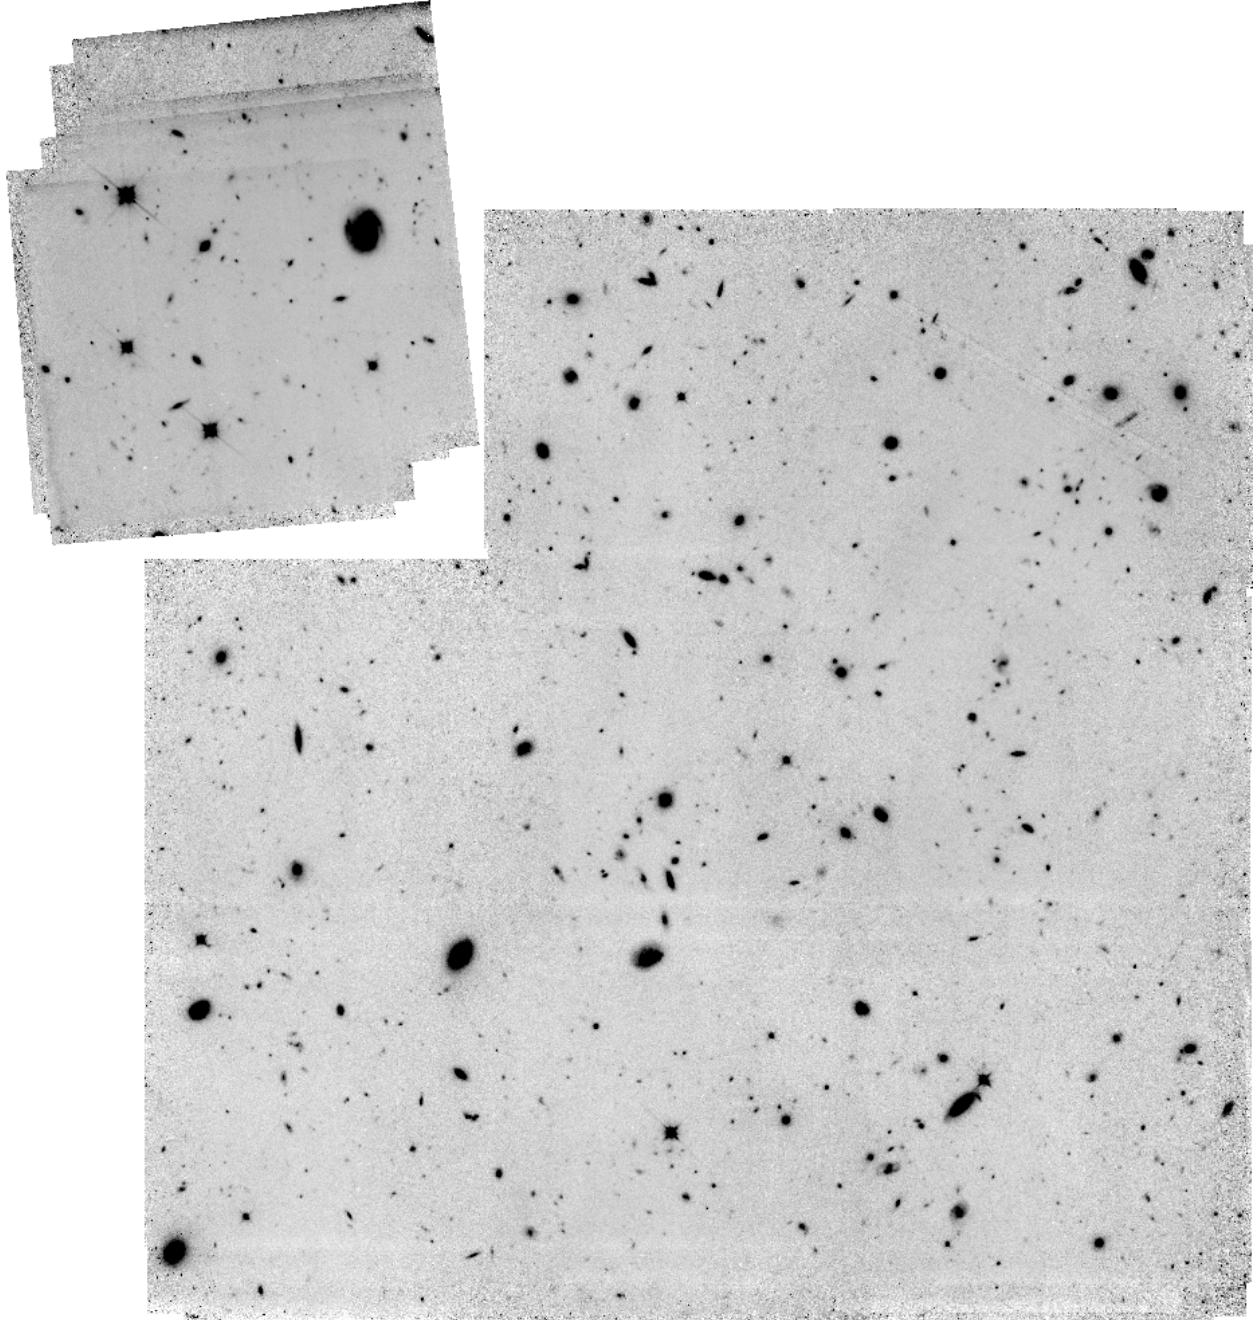


Fig. 1.— Least-squares calibrated *HST*/NICMOS F110W image of the HDF-S region (small field at upper left) and HDF-N region (large field). The model used in the least-squares solution includes detector gains and offsets, as well as variable offsets per detector quadrant (Eq. 2). The grey scale is linear and intentionally stretched to show low level calibration errors and the faintest detectable sources. Sky pixels are the same size as the $0.^{\prime\prime}2 \times 0.^{\prime\prime}2$ pixels of the NIC3 camera. For the HDF-S field N is upward; for the HDF-N field, N is rotated 122° counterclockwise from vertical.

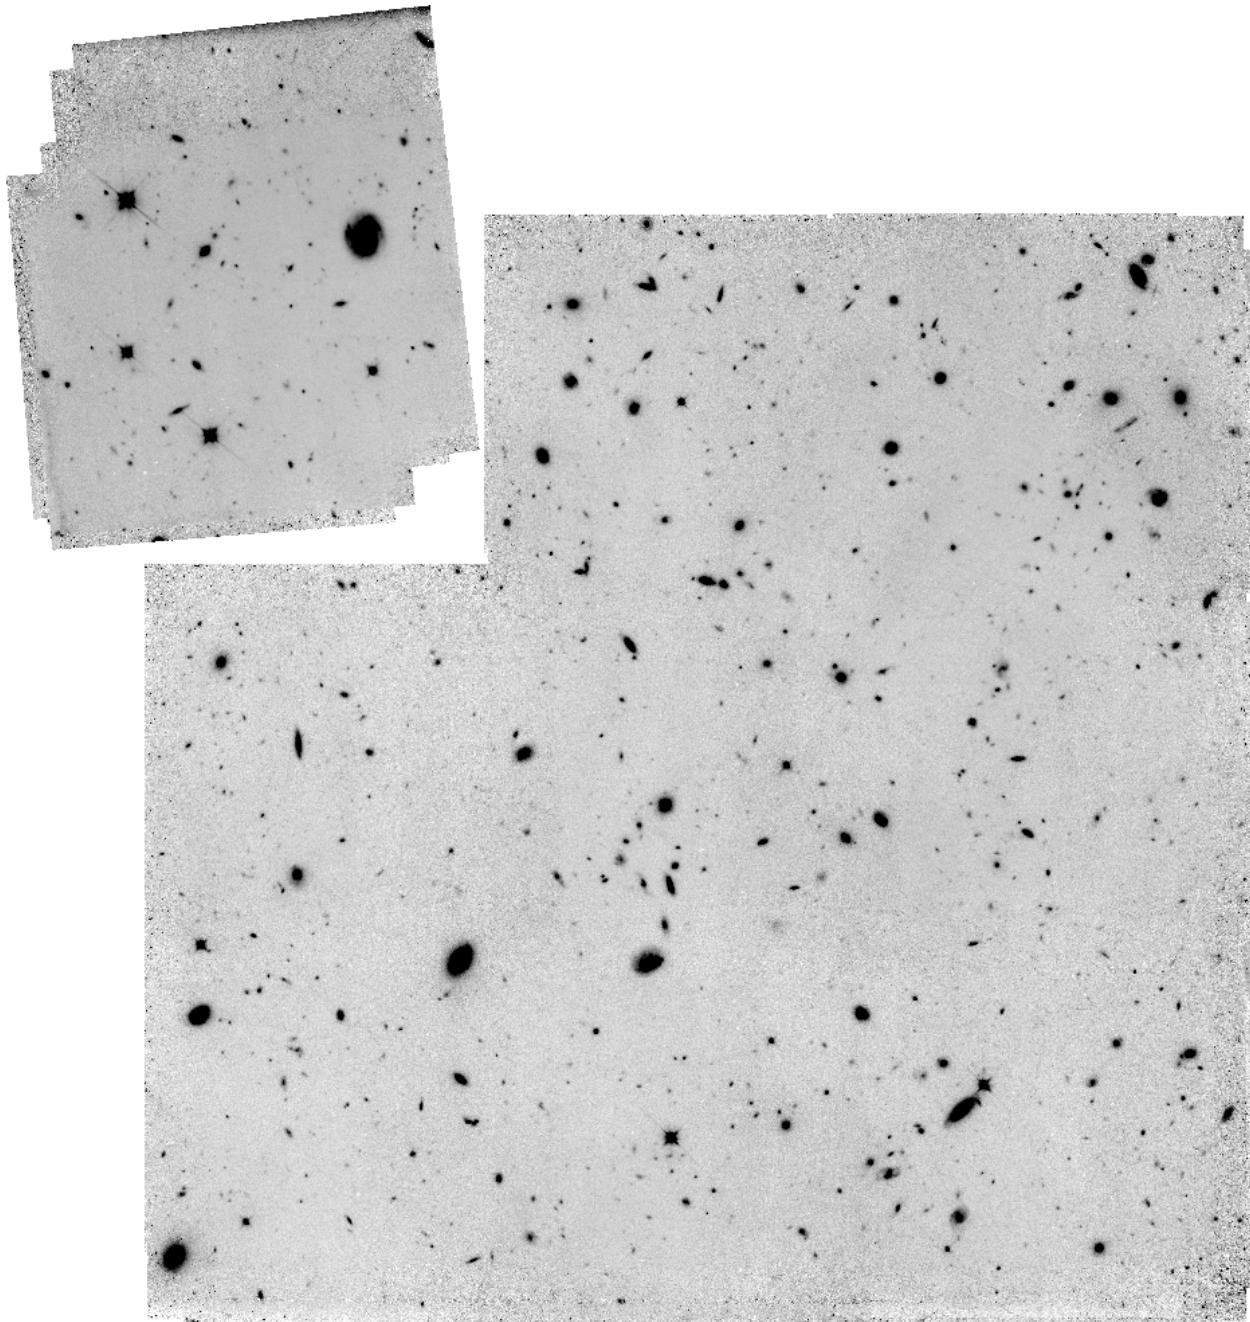


Fig. 2.— Same as Fig. 1, but with the vignetted bottom edge of the detector given lower weights in the final mosaicking of the calibrated individual frames. This reduces the artifacts caused by unmodeled effects in this portion of the images. Display range is identical to that of Fig. 1. As oriented here, the lower edge of the detector is mapped into the upper portions of the HDF-S region and the lower portions of the each of the 8 tiled subregions of the wide HDF-N region.

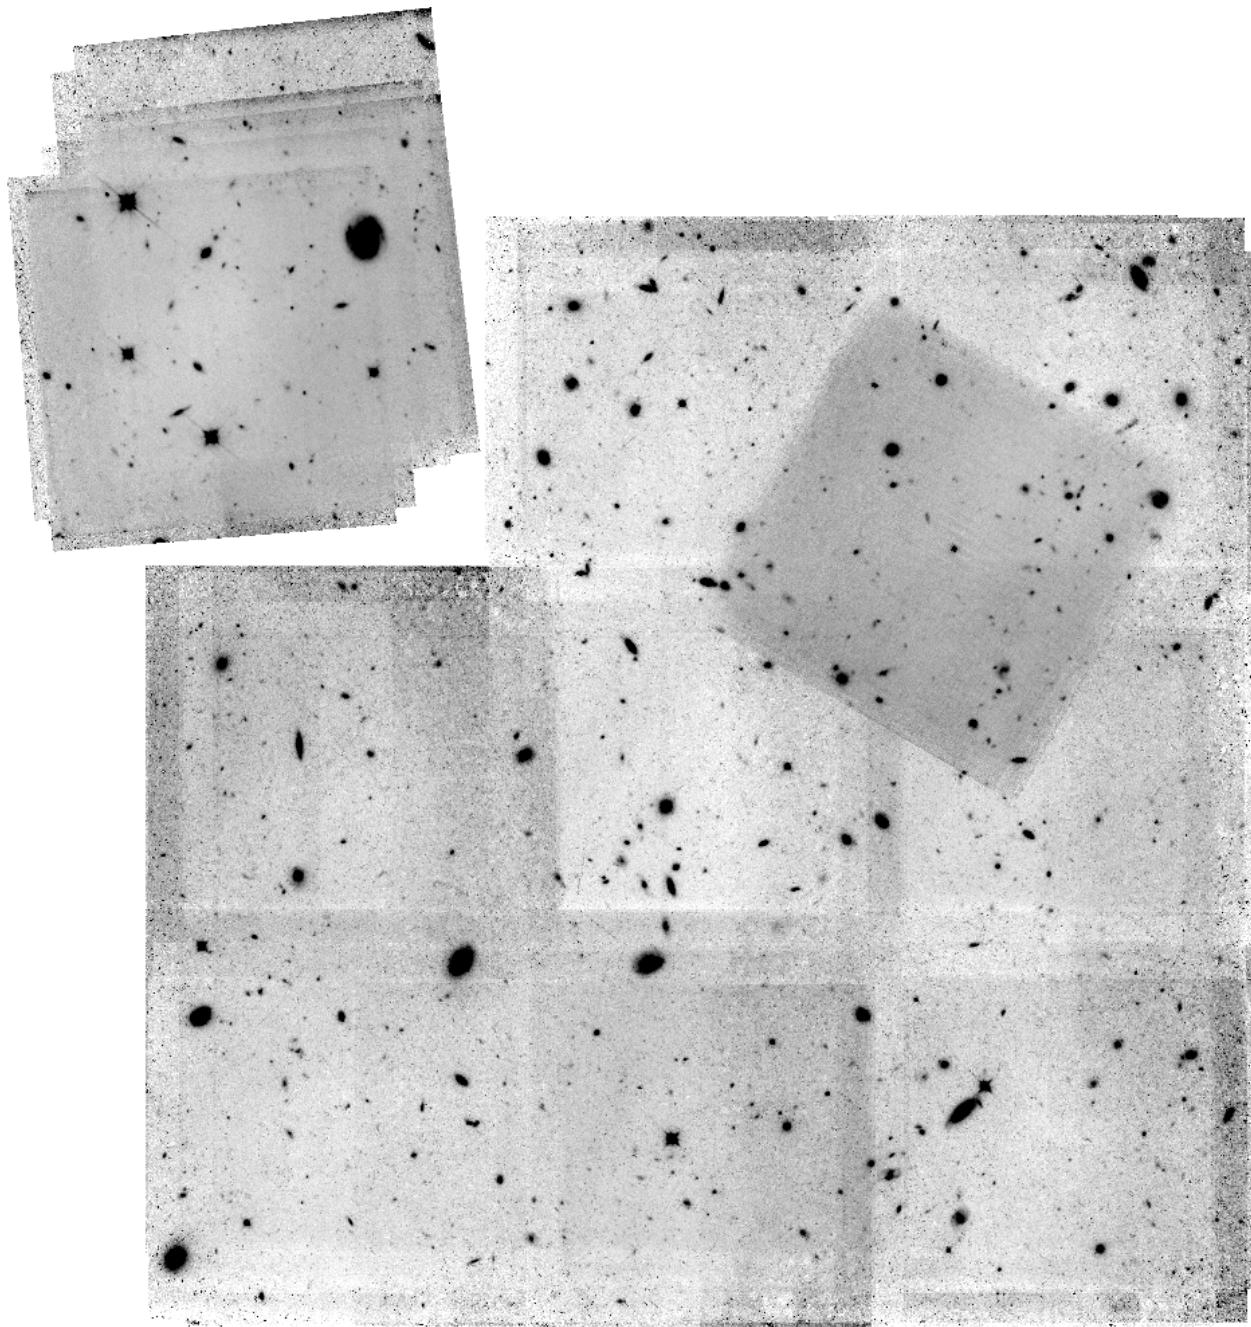


Fig. 3.— Same as Fig. 1, except here the model (Eq. 1) did not include a variable offset per detector quadrant (i.e. the “pedestal” effect). Display range is identical to that of Fig. 1.

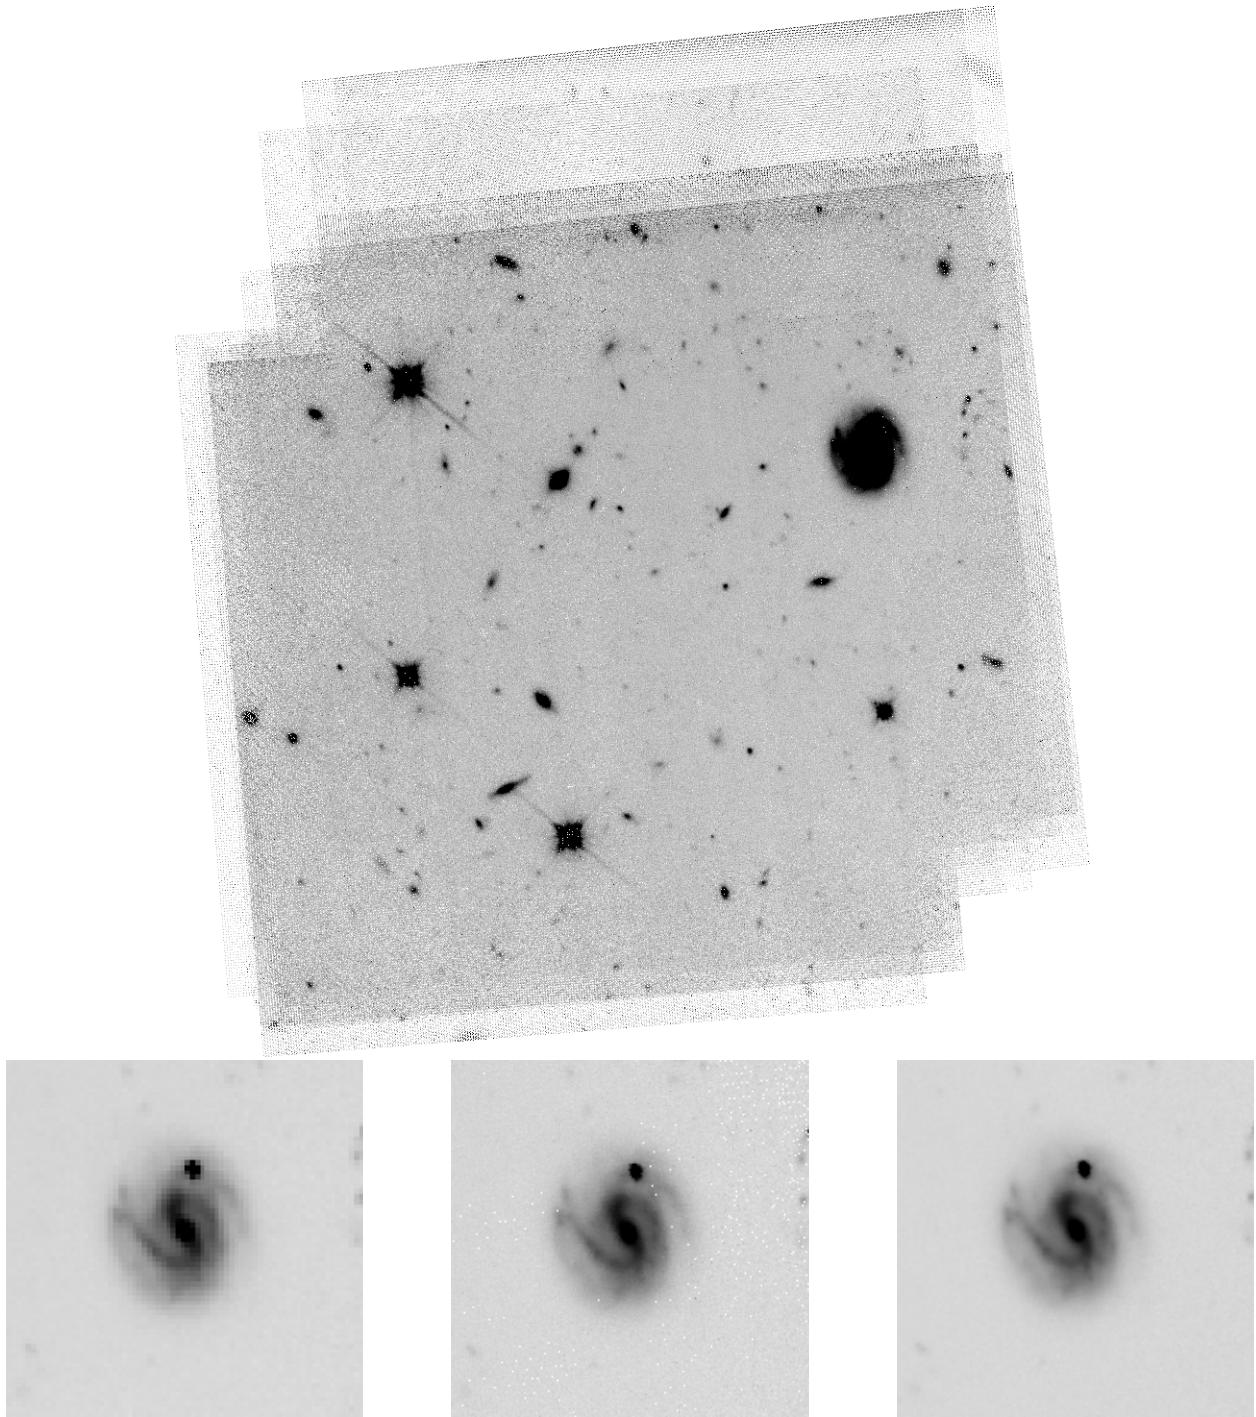


Fig. 4.— (a) The HDF-S region mapped on $0''.074 \times 0''.074$ pixels using the least-squares calibrated data from Figure 2. (b) Expanded view of the large bright galaxy in the HDF-S field (HDF-S J223249.56-603835.5). Left to right: Least-squares result on default pixel scale (from Fig. 1), least-squares result on fine pixel scale (from Fig. 4a), STScI calibrated and drizzled result.

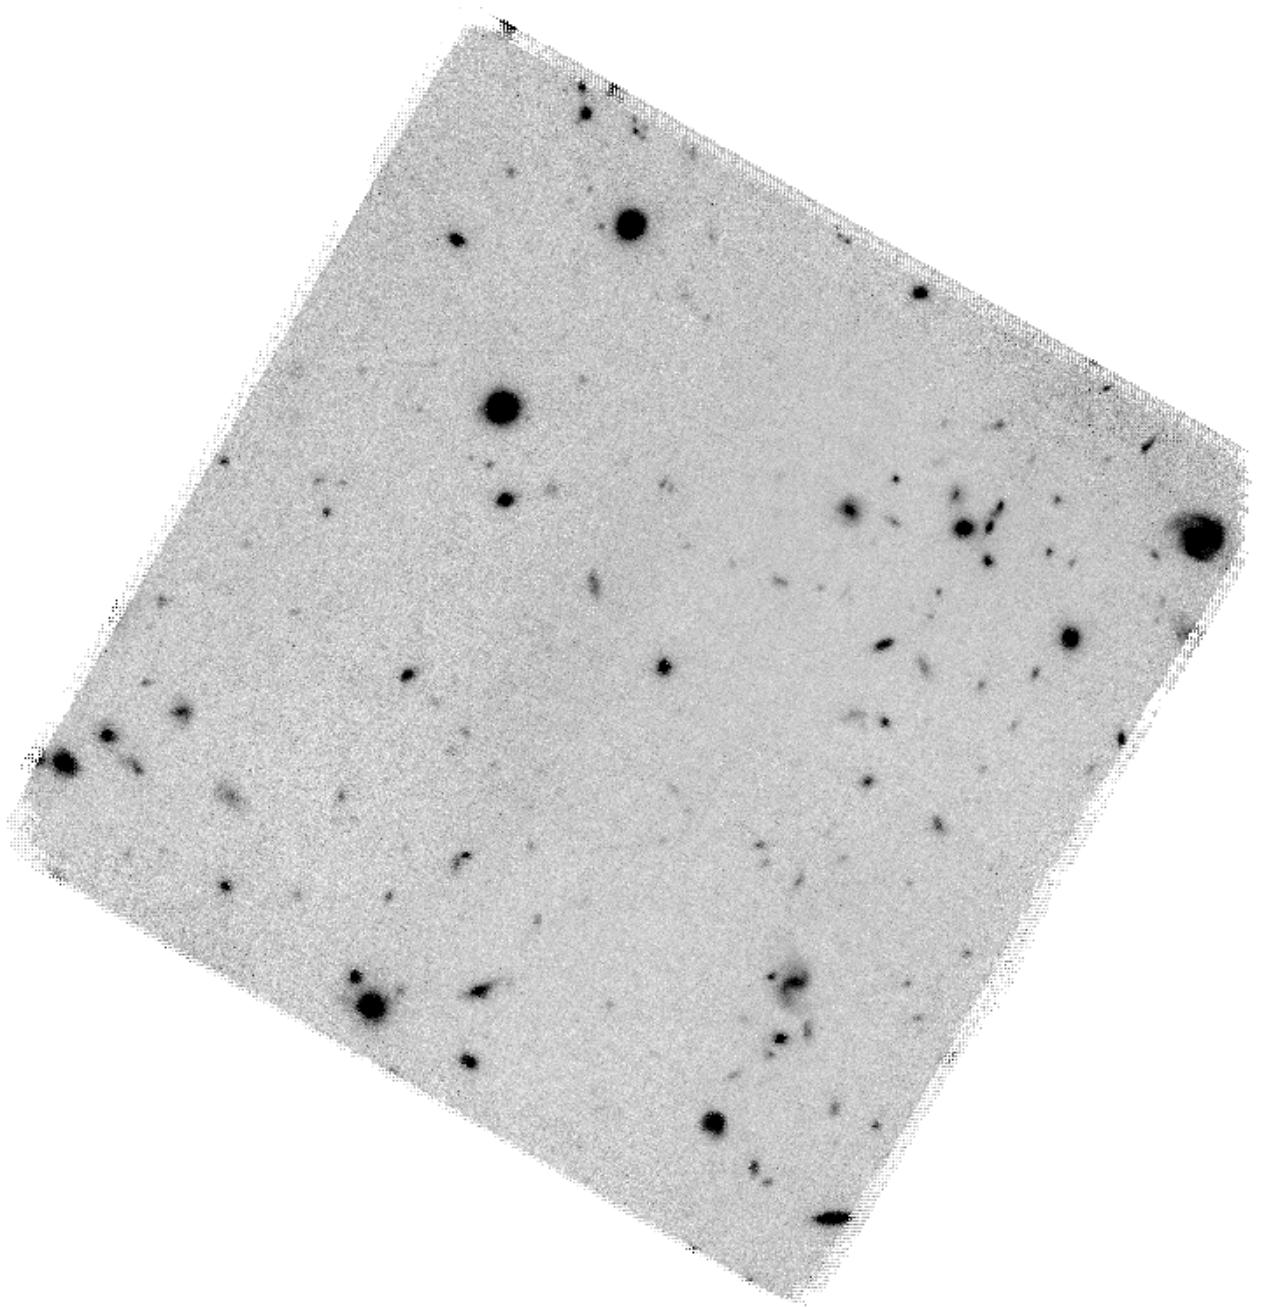


Fig. 5.— The HDF-N region of Thompson et al. (1999) mapped on $0''.1 \times 0''.1$ pixels using the least-squares calibrated data from Figure 2. Regions that are poorly covered at this fine pixel scale are blanked.

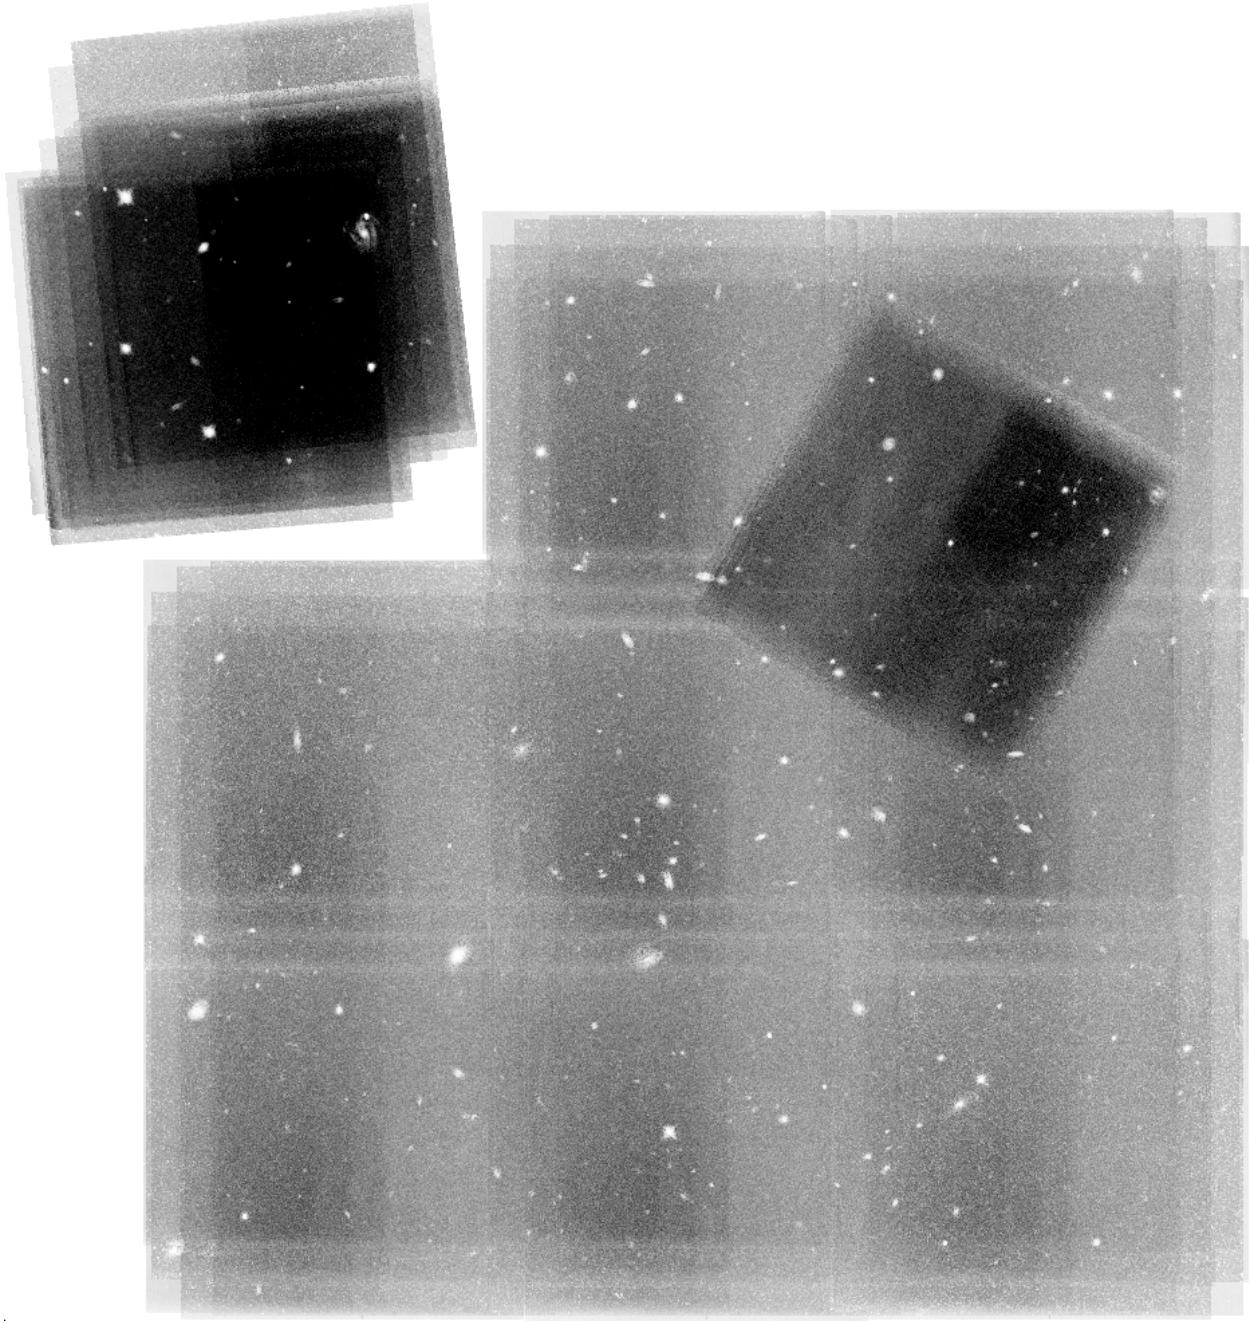


Fig. 6.— This image represents the weight of the observations. The square root of the weights, $\text{diag}(\Psi^{-1})^{0.5}$, is plotted with darker regions corresponding to higher weights. The number of integrations can be seen, as well as the more sensitive area of the detector. Bright sources appear with lower weights because of the additional noise due to photon statistics.

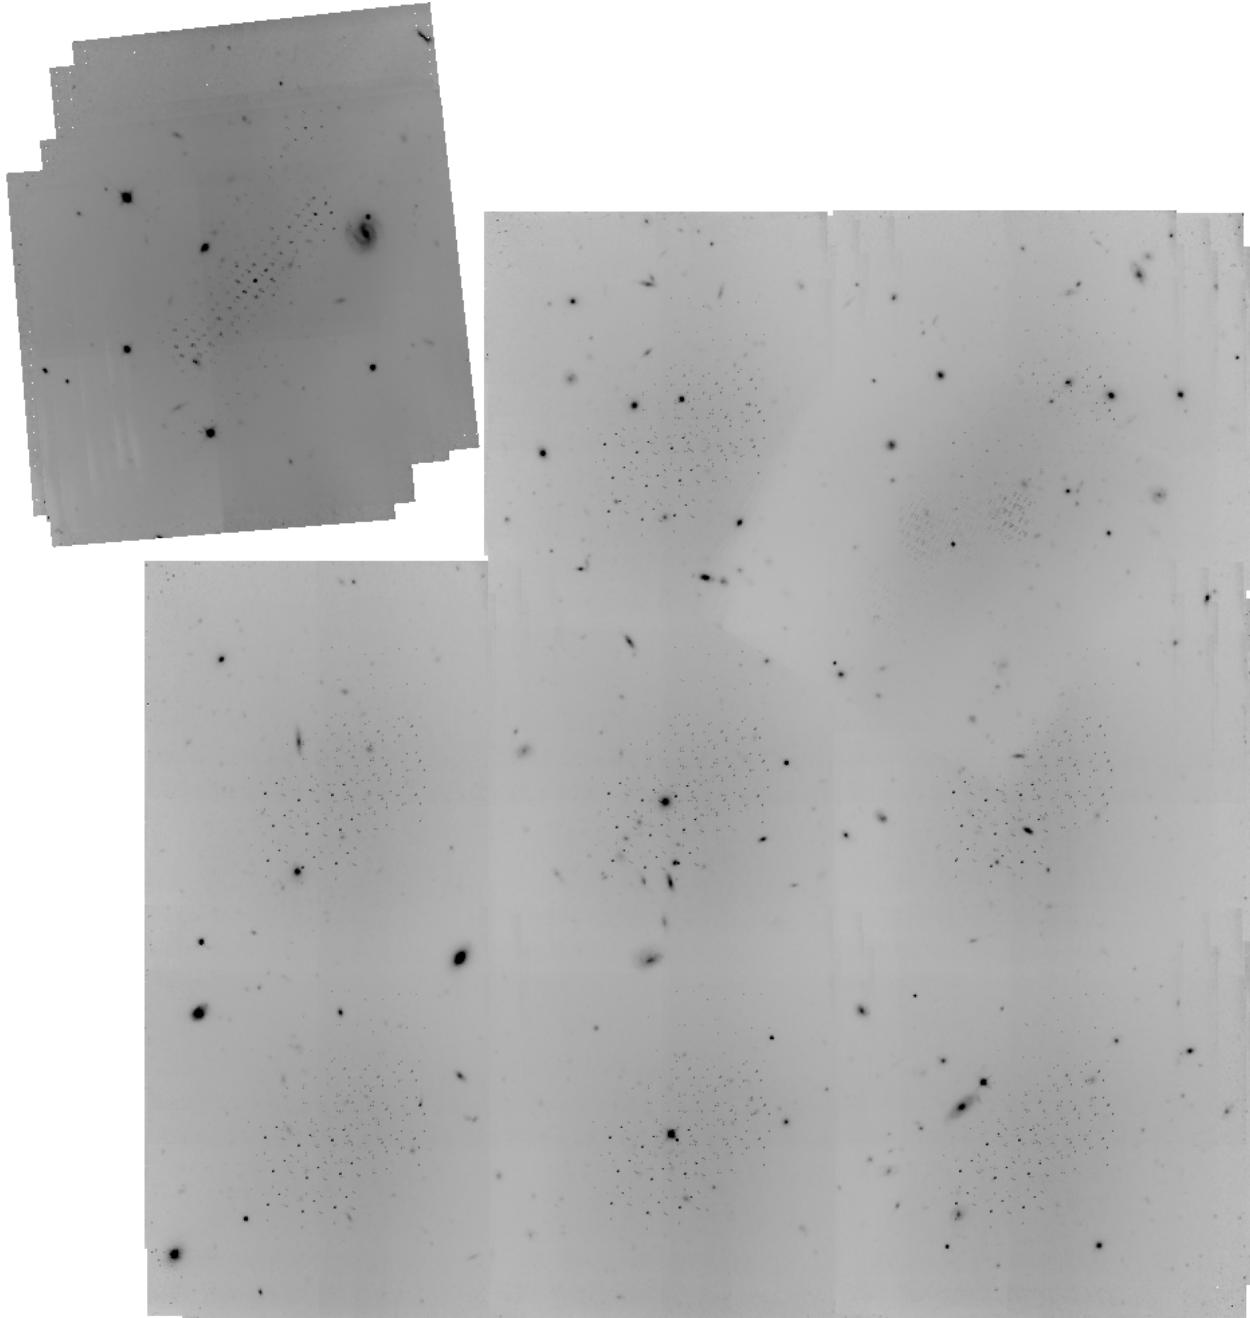


Fig. 7.— This image displays one row of the full covariance matrix, Ψ . The row corresponds to the pixel at the center of the pattern in the HDF-S region (upper left). The covariance is strongest between sky pixels that were observed with the same detector pixel, via the dither pattern. The dither pattern of the HDF-S region is visible in that region, and the 3×3 pattern used in the large area of the HDF-N is visible there. Bright features in the sky appear strongly in the covariance matrix because they are used in deriving the gain.

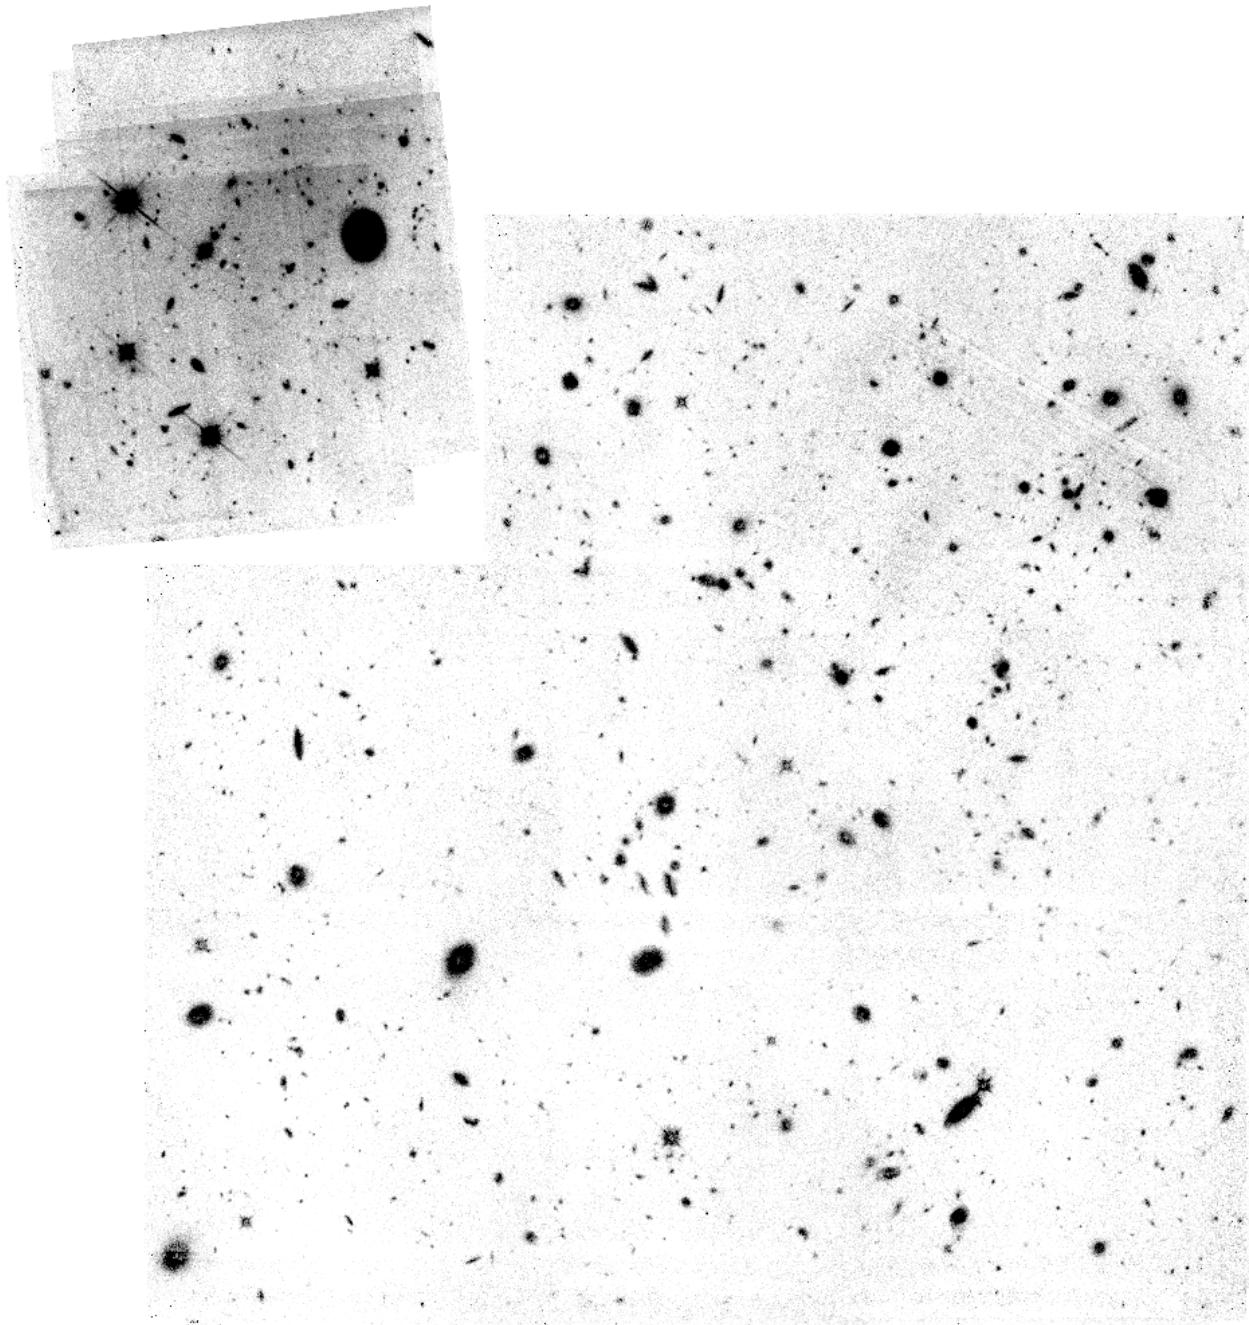


Fig. 8.— This image shows $[S^\alpha - \text{median}(S^\alpha)] * (\Psi^{-1})^{0.5}$, the significance (signal-to-noise ratio) of the final image. The display range is from 0.0 to 10.0 (white to black). Note the bright features more than overcome the additional noise from photon statistics (cf. Fig. 6).

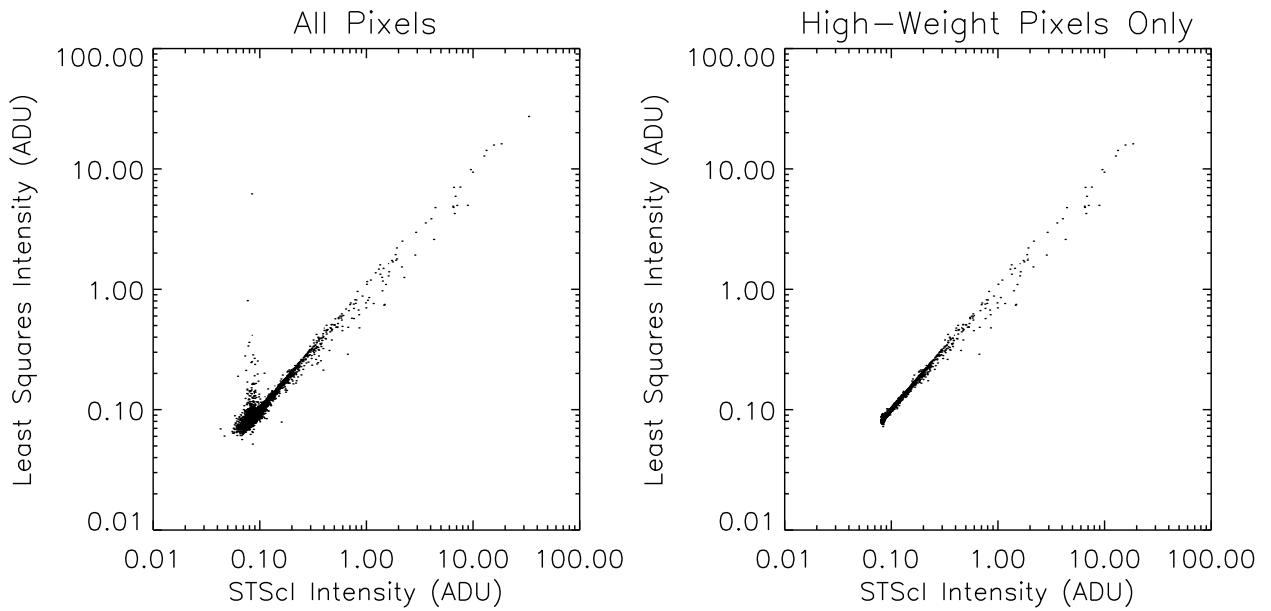


Fig. 9.— Correlation plots between the pixel intensities of the STScI produced HDF-S image and our least-squares calibrated image (Fig. 1). The left panel shows all pixels, while the right shows only those that are assigned high weights in the mosaic of the STScI calibrated images. This eliminates most of the residual cosmic ray effects.

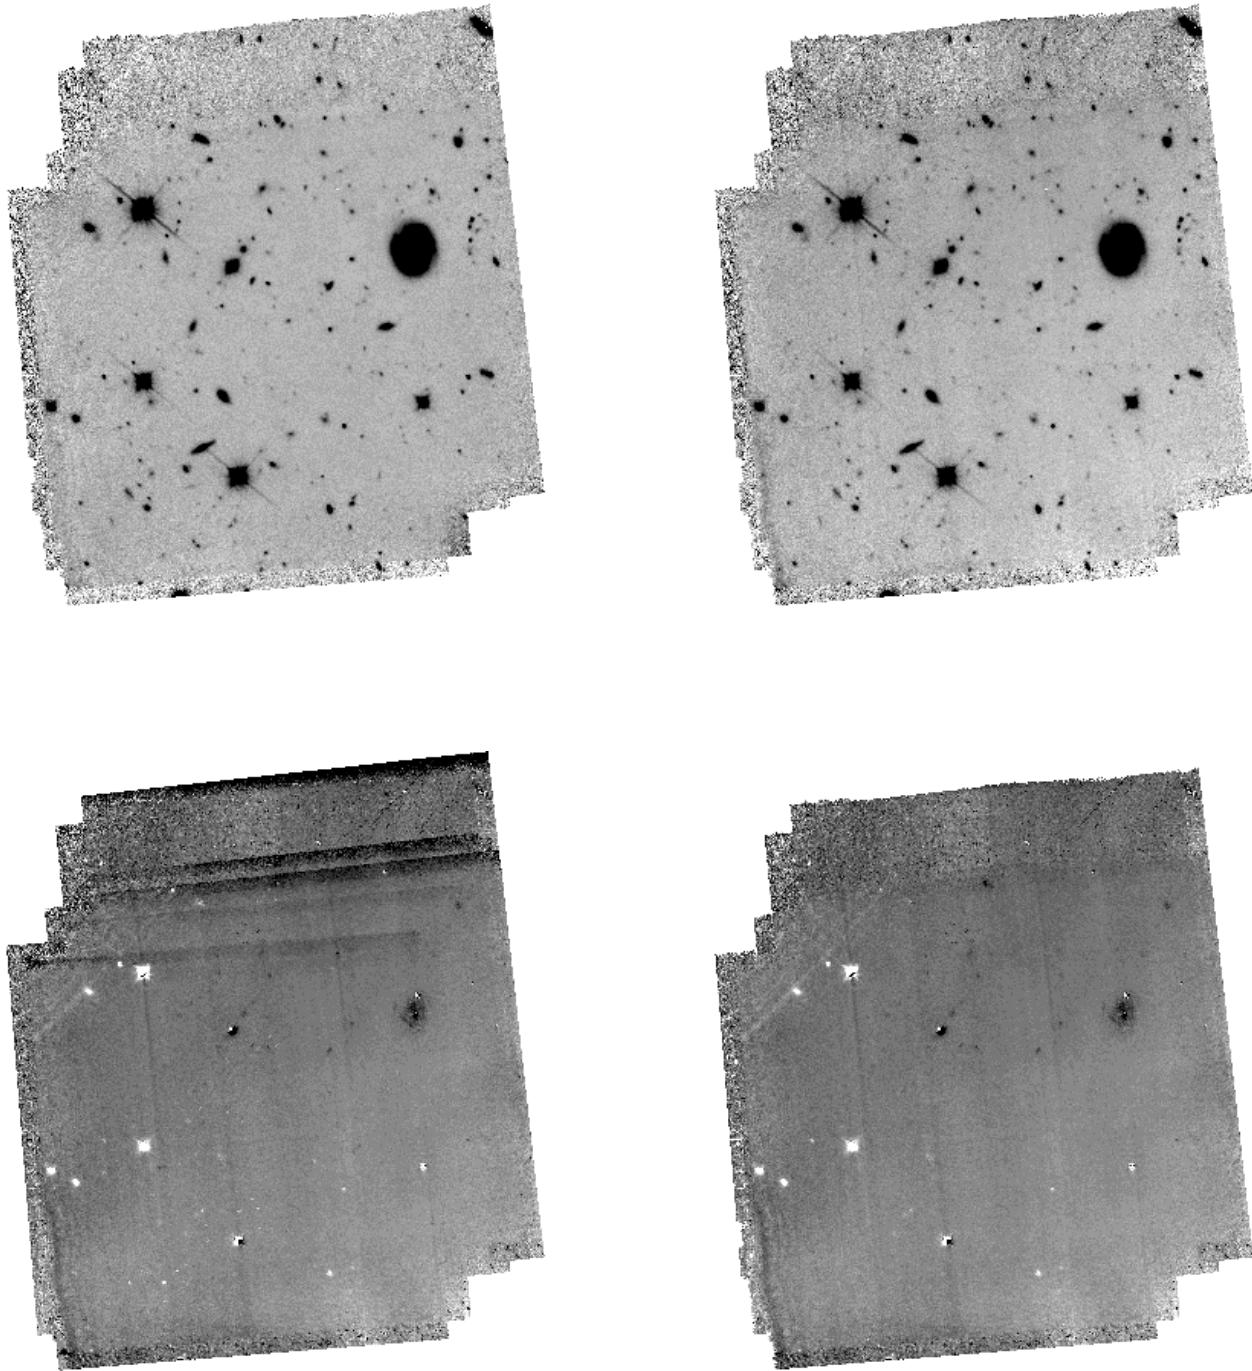


Fig. 10.— Left to right, top to bottom: The STScI calibrated HDF-S data mapped onto $0.^{\prime\prime}2$ pixels; the least-squares calibrated data, reweighted along the vignetted edge, and mapped into $0.^{\prime\prime}2$ pixels (Fig. 2); difference between the STScI image and the least squares calibrated image (Fig. 1); difference between the STScI image and the reweighted least-squares calibrated image (Fig. 2 and top right).

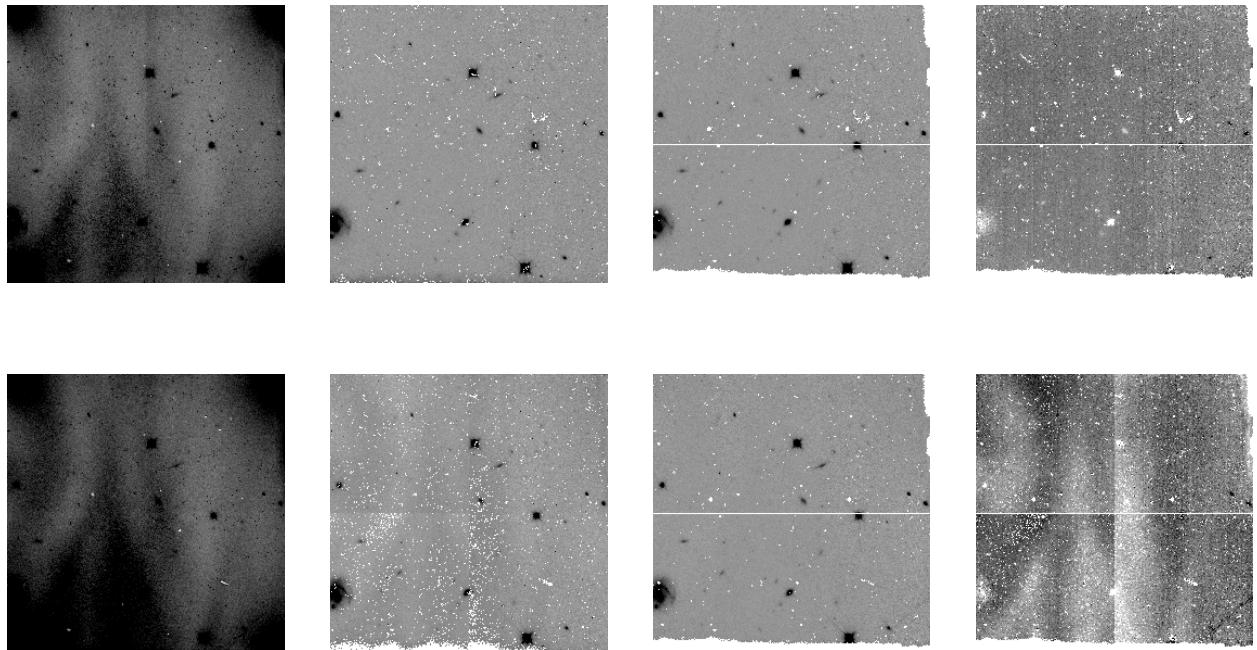


Fig. 11.— Top row shows: a single raw image; the image after the least squares calibration is applied; the STScI calibrated version of the image; the difference between the least-squares and STScI calibrated versions. The calibrated images (and difference) are masked at bad pixels as designated by the respective procedures. The second row shows the comparable steps for a single frame for which the least-squares calibration is not so successful. Defects in the form of residual gain or quadrant errors can be seen in many of the individual least-squares calibrated frames, but they are generally much weaker than those shown in the second row.

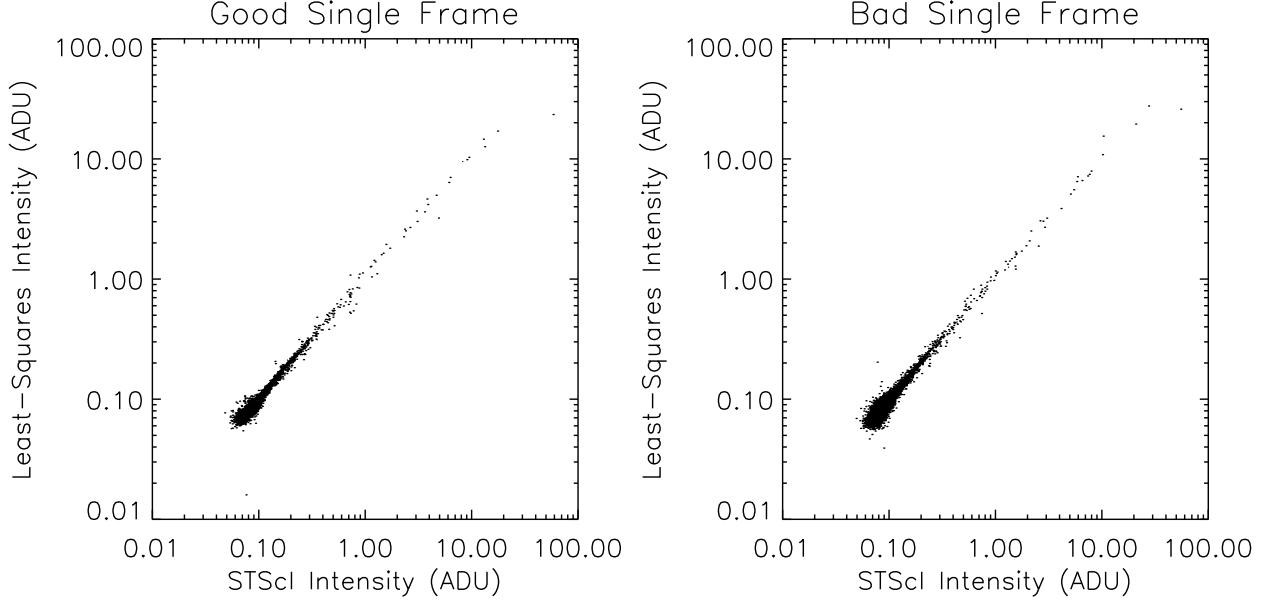


Fig. 12.— The pixel to pixel correlations of the least-squares and STScI calibrated frames from Fig. 11 are shown.

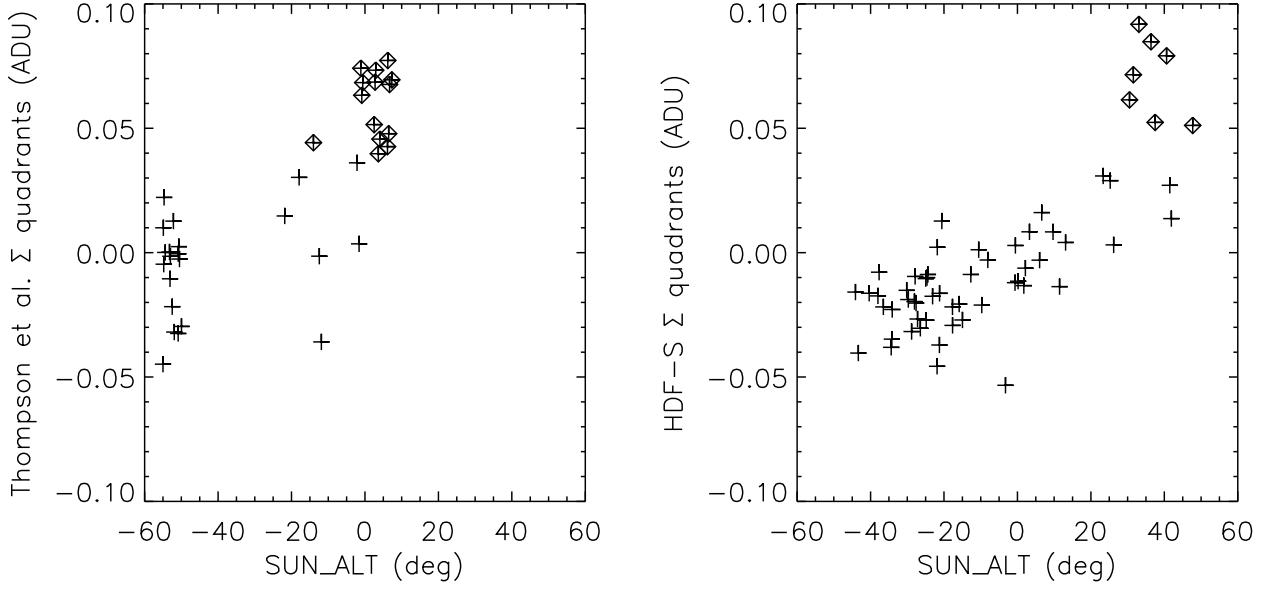


Fig. 13.— The sums of the derived quadrant offsets (pedestals) for each frame used in the Thompson et al. HDF-N observations (left) and the HDF-S observations (right) are plotted as a function of the elevation of the Sun above the Earth limb as indicated by the SUN_ALT keyword in the FITS headers. All frames are marked with “+” symbols. Frames where the sum of the quadrants offset, F^q , is > 0.037 are overplotted with “◊” symbols.

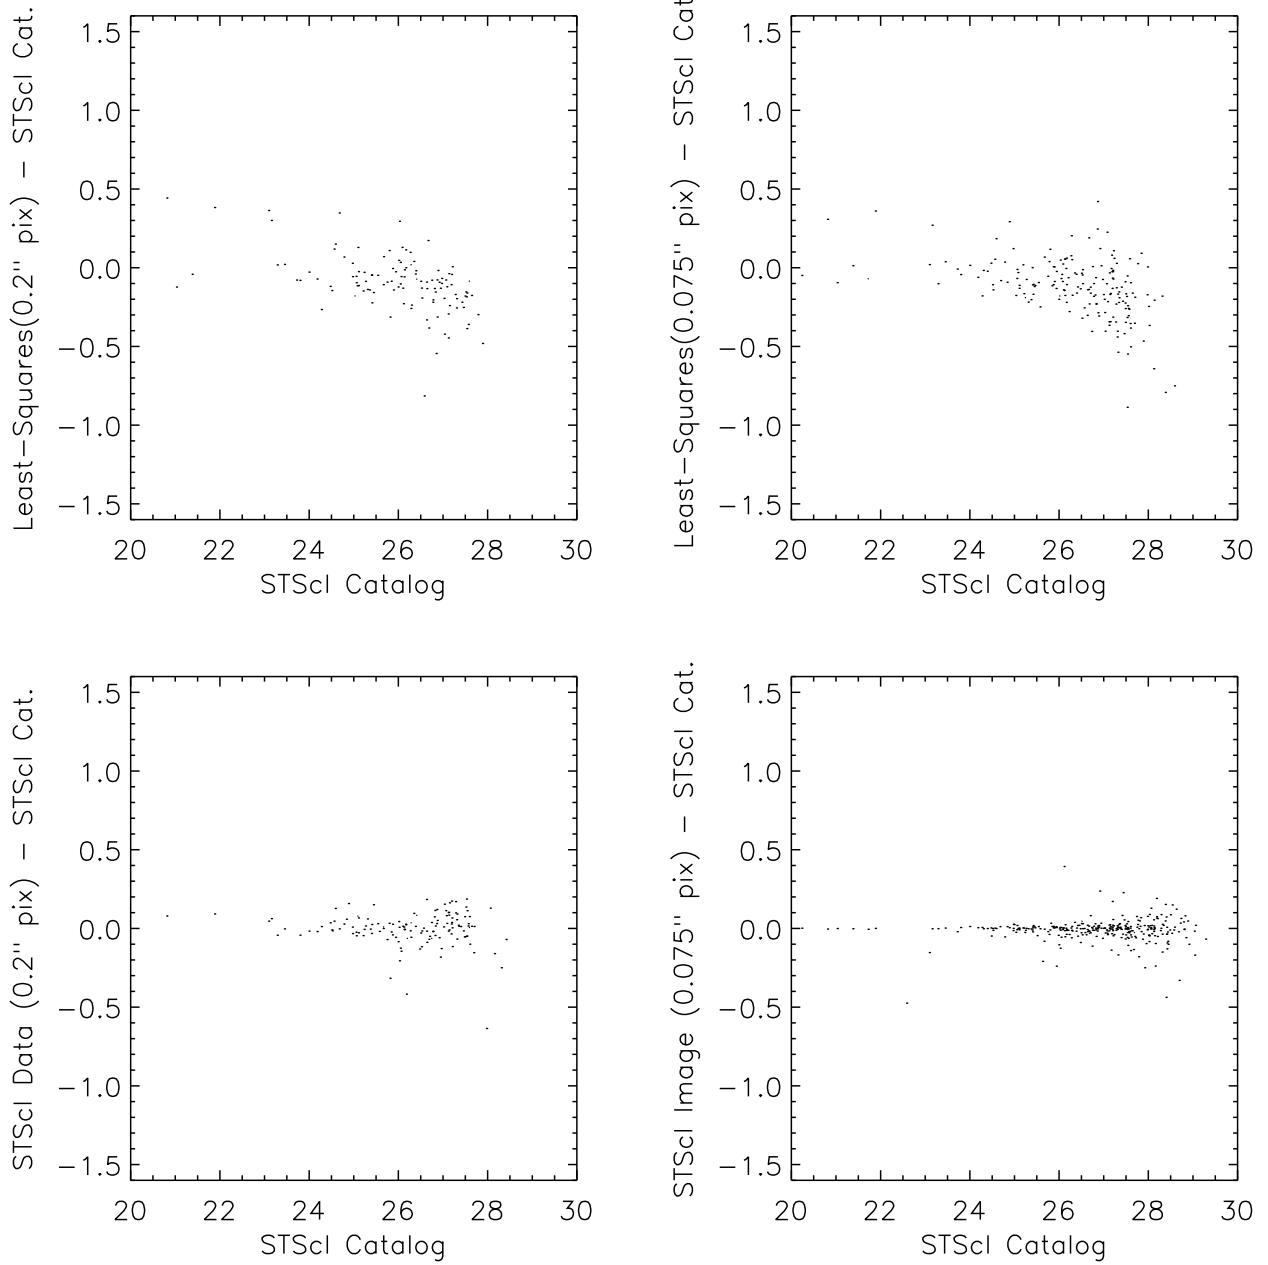


Fig. 14.— HDF-S magnitude differences (Δm) between SExtractor results and the STScI catalog ($0''.5$ apertures) for four different images: least-squares calibrated data on $0''.2$ pixels [$\sigma(\Delta m) = 0.19$ mag], least-squares calibrated data on $0''.075$ pixels [$\sigma(\Delta m) = 0.19$ mag], STScI calibrated data on $0''.2$ pixels [$\sigma(\Delta m) = 0.11$ mag], and the STScI calibrated and drizzled map on $0''.075$ pixels [$\sigma(\Delta m) = 0.07$ mag].

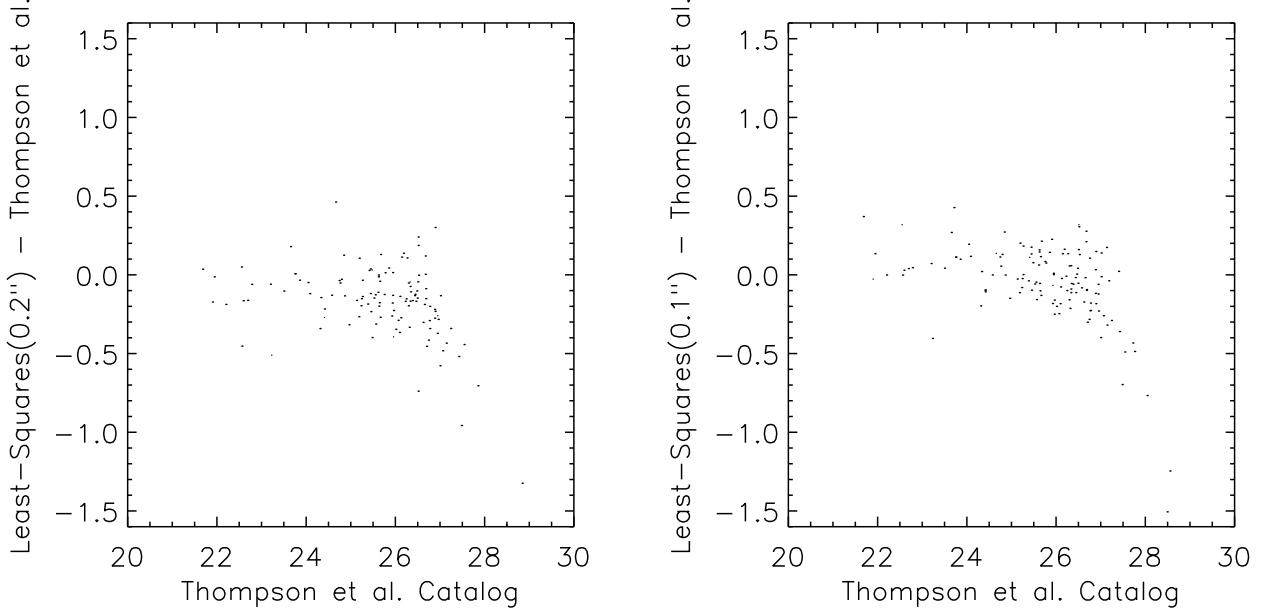


Fig. 15.— HDF-N magnitude differences between SExtractor results and the Thompson et al. (1999) catalog ($0''.6$ apertures) for two different images: least-squares calibrated data on $0''.2$ pixels, least-squares calibrated data on $0''.1$ pixels.

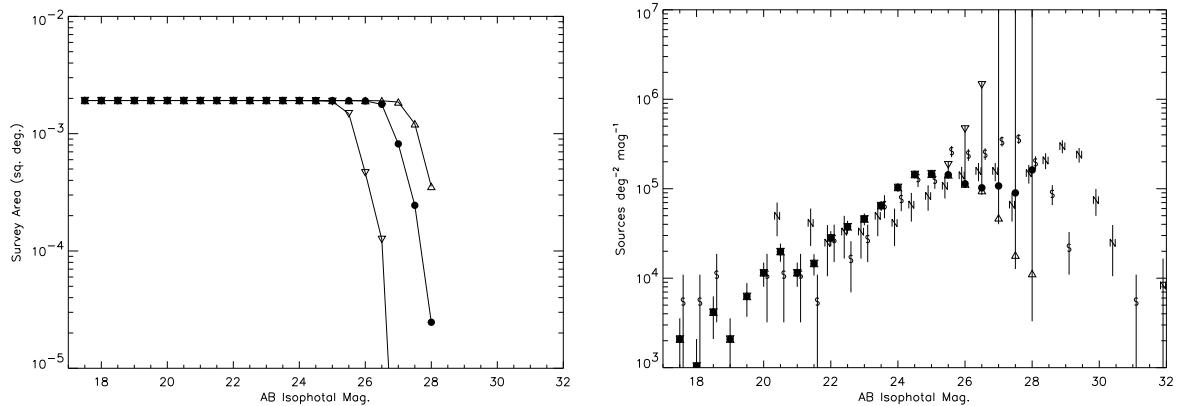


Fig. 16.— (a) Effective survey area as a function of AB magnitude (F110W) for the combined data sets. The effective areas for sources that are 3.26 , 6.03 , and 18.45 pix^2 in size are shown with the symbols Δ , \bullet , and ∇ respectively. (b) Differential F110W source counts for the combined least-squares calibrated image (Fig. 2) assuming survey areas as depicted in (a) (symbols: Δ , \bullet , ∇), the Thompson et al. (1999) HDF-N catalog (symbol: N), and the STScI HDF-S catalog (symbol: S). Error bars show $\pm 1-\sigma$ Poisson uncertainties for each bin, and the range of correction for survey area as a function of source size. Bins for the referenced HDF-S and HDF-N data are displaced by ± 0.1 mag for clarity.