SPACE TELESCOPE SCIENCE INSTITUTE
Operated for NASA by AURA

Instrument Science Report ISR WFC3-2011-11

Sky Flats: Generating Improved WFC3 IR Flat-fields

N. Pirzkal, J. Mack, T. Dahlen, E. Sabbi
May 2, 2011

ABSTRACT

A significantly improved set of flat-fields are now available and are currently used as part of the WFC3 calibration pipeline. We describe the creation and testing of new in-orbit flat-field corrections for the WFC3 IR channel. While high signal to noise ground based flat-fields were generated prior to launch, photometry of dithered stellar fields showed that these flat-fields failed to fully flatten the large scale structure of the WFC3 IR flat-fields. In this ISR we show how we generated a correction to the ground based flat-fields using thousands of IR observations. This correction, or sky delta flat-field (SD-flat in this ISR), appears to be both wavelength and time independent and is stable down to better than 1% over most of the detector. Photometric accuracy using new corrected flat-fields is better than 0.5% (peak to peak variation of -1.5/+1.6%) if one avoid being within 128 pixels of the edge of the detector. For the "wagon-wheel" region and the edge of the detector, photometric accuracy is reduced to about 0.8% (peak to peak variation of -2.0/+1.9%).

Introduction

A collection of flat-fields (hereafter referred to as CASTLE LP-flats, and generated for the IR channel of WFC3 prior to launch) were created for all WFC3 filters using the HST high fidelity simulator CASTLE to illuminate the IR channel. The document WFC3-ISR-2008-28 fully describes these measurements. The IR WFC3 CASTLE LP-flats were intended to provide solid small spatial scale, or pixel-to-pixel, flat-fielding information in order to be combined with in-orbit flats obtained using the WFC3 internal lamp.

Stellar photometry in dithered observations of the star clusters Omega-Cen and 47 Tuc, obtained during the Servicing Mission Observatory Verification (SMOV) proposal CAL-11453, indicated that the WFC3 IR CASTLE LP-flats did not fully correct for large scale structure in WFC3 IR data These results were confirmed during the Cycle 17 calibration program CAL-11928.

In an ideal situation, one would like to routinely generate accurate in-orbits flats by exposing the full telescope to a uniform external light source. One approach to achieving this is to take advantage of the fact that the near infrared sky background is high enough to provide a uniform source of light in long exposures, if one carefully avoids astronomical sources. This is however time consuming because the NIR sky levels are relatively low (~0.3 to 1.0 e-/s/pixel) and therefore require deep imaging. It is moreover difficult to implement since there are no empty field one could point the telescope for extended periods of time. One method is to point the telescope at the bright Earth limb, which should be relatively featureless (proposal CAL-11917).

An alternative to this approach, is to use the WFC3 IR data accumulated since the successful installation of WFC3 on HST. While the available data are not devoid of astronomical sources and the NIR background does vary both as a function of time and telescope pointing, we devised a method to generate sky delta flats (SD-flats in this ISR) by combining together as many IR images as possible while masking out sources in the field and normalizing the overall sky background level. Details of the method used are given in the next section. In the rest of this ISR we describe the SD-flats, estimate their accuracy, and constrain their wavelength dependence. We show how applying a grey, multi-filters correction to all ground based flat-fields (which we refer to as CASTLE LP-flats thereafter for clarity) allows for a significant improvement of WFC3 IR flat-fielding. We finally show how the SD-flats compared to L-flats (which are flats that describe only the large scale flat-field dependence of the instrument) derived using photometry only, as well as to newer bright Earth limb flats that have been under construction since the release of our new IR flats.

Flat-field Nomenclature

HST data calibration involves the creation of different types of flat-fields, or flats. We summarize some of them here for the context of this ISR. Historically, HST flats have been created using two different types of flats: the LP-flats which contain information about the pixel to pixel variations (and are derived from very high signal to noise measurements on the ground) as well as large scale variations, while the L-flats only contain information about the larger scale variation of the entire optical system. One problem with LP-flats obtained from the ground is that the CASTLE HST simulator is not a perfect replica of the HST optical path and the resulting flats cannot be expected to perfectly flat-field observations taken with HST. Slight illumination difference between the ground test equipment and HST can result in errors of up to several percents. New L-flats corrections are commonly derived once an instrument is installed on HST using a variety of methods. These are then used to create new LP-flats that correct for flat-fielding on both small (pixel to pixel) and larger scales. LP-flats are what are used by users and the WFC3

pipeline to flat-field data. For clarity, Table 1 summarizes the different types of flats discussed in this ISR.

Flat-field Type	Description
L-flat	Low-frequency correction to detector sensitivity, due to differences in-flight versus ground calibration computed in one of three ways: from 'SD flats', from 'Earth flats', or 'Stellar L-flats'
LP-flat	Detector response correction image, including both pixel-to-pixel sensitivity and low-frequency modulations
CASTLE LP-flat	Also commonly referred as a "ground flat". An LP-flat is obtained from uniform illumination via the CASTLE simulator during Thermal Vacuum 3 ground testing. We refer to these as CASTLE LP-flats in this ISR as this is a more accurate description
SD-flat	Sky Delta flat, which is a correction to an LP-flat and is created by combining background sky observations
Earth flat	LP-flat obtained from observations of the bright Earth limb (may be sunlit or moonlit)
Stellar L-flat	L-flat derived by moving stars to different regions on the detector and measuring changes in response
Sky flat	LP-flat obtained from observations of the sky as seen by HST. Can be created by correcting an existing LP-flat by an SD-flat, as is discussed in this ISR

Table 1: This table lists the different types of flat-fields discussed in this ISR and briefly summarizes what they are and how they are connected to one another.

Building the IR SD-flats

We computed corrections (SD-flats) to the existing ground based WFC3 IR flat-fields by combining calibration and GO IR data. These data (FLT files) were already calibrated by the HST archive pipeline and hence already flat-fielded by the existing CASTLE LP-flats. We used flat-fielded data in order to start this process with data that were already close to flat and for which we could derive accurate background values. The structure and amplitude of the signal present in the SD-flats are direct measurements of the inaccuracy present in the ground flat-fields. Had the ground flat-fields been perfect, the SD-flats we derive would be completely spatially flat, with a mean value of unity, no spatial variation, and with a noise level expected from Poisson statistics.

We started by first identifying observations of astronomical fields that are relatively devoid of large extended sources. The proposal IDs that we used are shown in Table 2. We restricted ourselves to using datasets with integration times larger than 300 seconds to ensure the presence of an adequate amount of sky light and that the data are not read noise limited. At approximately 0.5e-/s/pixel, a 300s exposure should contain about 150e-/pixel from the background sky, with a signal-to-noise larger than 10.

With each dataset, we used Sextractor, in an aggressive detection mode (DETECT_MINAREA=4, DETECT_THRESHOLD=0.85) to create a segmentation map of the image. This map is an object map generated by Sextractor where pixels that are part of a detected object are set to a non zero value. We generated an inverse object mask from these by setting all non zero values to 1 and leaving all empty pixels to 0. We then "grew" this mask by convolving it with a gaussian with a full width at half maximum of 10σ pixels. As a final step we then set pixels with values larger than 0.03 in this convolved mask were set to a value of 1 while all other pixels were set to 0. This effectively grows the size of a point source to about 40 pixels, which we found to be necessary to remove contributions from the very faint edges of objects. We determined the appropriate size of the gaussian and final threshold value empirically and found they were a good compromise between masking out most pixels affected by astronomical sources (especially on the edges of these sources) and not masking out too many pixels. Typically, between 50% to 80% of all pixels in a given dataset were masked out. While this is a significant fraction of the available pixels, the random distribution of sources in the field combined with a large number of independent datasets ensured an equal coverage of the entire array.

In addition to excluding objects in the field, we also excluded pixels likely affected by persistence. This effect is discussed in WFC3 ISR 2010-17. We avoided such pixels in each dataset by examining all of the available data acquired using the WFC3 IR channel. We simply identified any pixel that was filled by more than 30000e- in any exposure taken within three days prior and added those pixels to our pixel mask.

Next, we estimated the amount of background in each individually masked image. The background, which varies drastically from one image to the next, needs to be normalized out of each dataset. This is a necessary step before we could combine these data to generate an SD-flat. We estimated the background level by computing the median of each image, ignoring the contribution of masked out pixels. We then normalized the masked exposure by this median background

level. Note that the median value is a good estimate of the background level in this particular case because the input data were already flat-fielded using the LP-flats and were thus flat at the level of a few percents. We only used datasets in which the mode and median of the background were found to be within 5% of one another. This extra step allowed us to automatically avoid most datasets with residual sources or scattered light.

At this stage, we had a series of masked, normalized images which were combined together, as a final step, again ignoring the value of any masked pixel. In summary, for each pixel (i,j) we generated a list of background values from the individual normalized and masked images. Starting with N input datasets, we ignored all the M masked values (which by definition were contaminated by the light of some astronomical source) and computed the median, per pixel, of the remaining (N-M) values. Using a number of input images N that was large enough (N on the order of several 100's) we found that we had enough N-M values left to compute a meaningful median at every position (i,j) on the detector. It should be noted however that this was only the case when N was large and not with less often used filters (i.e. F105W, F110W, F140W) for which the resulting SD-flat is more noisy.

Figure 1 illustrates this process using a single input dataset, the Sextractor initial segmentation map, the mask after being "grown" using a gaussian (σ=10 pixel) convolution, and finally the resulting masked and normalized image. We combined several hundred images like this one to generate in-orbit SD-flats for several medium and broad band IR WFC3 filters.

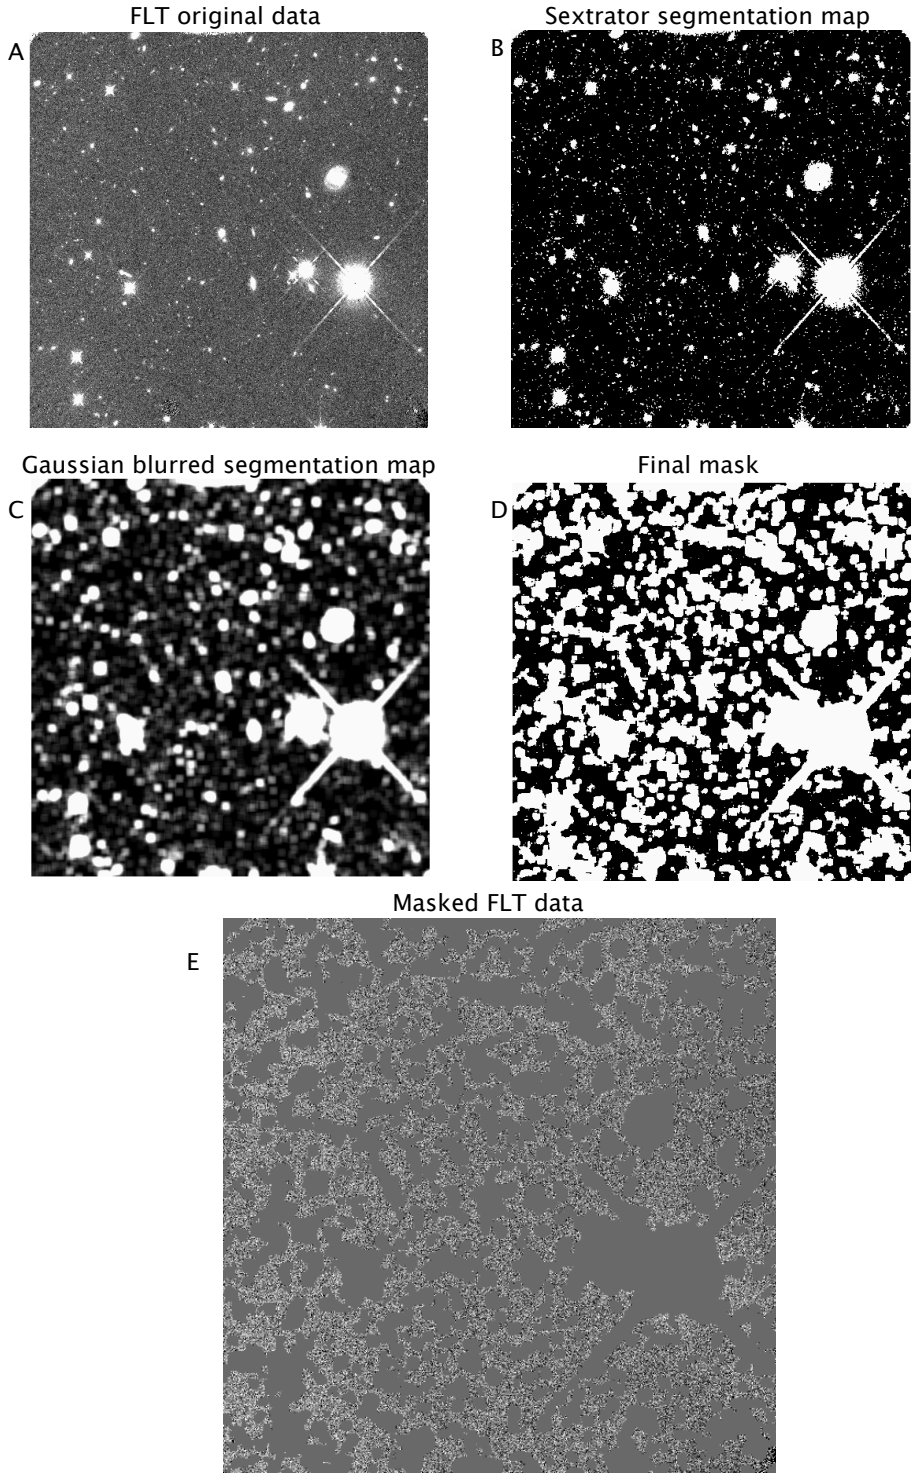

Figure 1: Mask creation example. Starting from the original FLT data (A), a Sextractor segmentation map is created (B). The latter is then gaussian blurred (C) and a mask is created from the result (D). The final masked image is shown in (E) and this masked image is then normalized by its median value.

SD-Flats: Sky Delta Flats

Using data taken between 09-2009 to 12-2010, we were able to generate SD-flats for F098M, F105W, F110W, F125W, F140W and F160W. Table 3 lists the number of datasets that were used to generated these flat-fields. The highest signal to noise flat field was generated for the F160W filter for which nearly one thousand individual FLT images were combined. As a result of the object masking and aggressive pixel rejection, only a subset of these datasets were used at any given pixel. Table 3 lists the maximum number of value used to compute the SD-flat value for each filter. While enough data existed to also generate (noisier) SD-flats in F098M and F125W filters, there were not enough data for F105W, F110W, and F140W: only about 100 datasets could be combined, resulting in SD-flat with considerable lower signal to noise per pixel (see column five in Table 3) as well as "holes", or regions on the detector where we had no data to combine. The average standard deviation shown in Table 3 is the mean, over the whole detector, of all the individual pixel standard deviations. The latter were computed from the list of N-M values for that pixel when combining the N datasets, as we described above. The average number of datasets shown in Table 3 is the number of datasets that were combined to produce any given pixel flat field value. Finally, the average signal-to-noise per pixel is the average total number of e- divided by the square root of that number that we used to compute the median, average and standard deviation discussed here. The latter is an upper limit to the signal to noise in each listed SD-flat since it does not include read-noise, dark-current nor any systematic errors.

In addition to individual IR filters, we also generated a grey SD-flat using data from all 6 filters, which we labeled as ALL/Grey in Table 3. The grey SD-flat was generated by combining nearly 2000 individual observations and not by averaging the other available SD-flats shown in Table 3. We estimate the signal to noise of this grey SD-flat to be twice of the F160W SD-flat signal to noise.

Proposal ID					
11108	11142	11149	11153	11166	11189
11202	11208	11343	11359	11519	11520
11534	11541	11557	11563	11584	11587
11597	11600	11644	11650	11666	11669
11694	11700	11702	11735	11838	11840

Table 2: List of proposals used to generate the SD-flats discussed in this ISR.

Filter	Maximum Number of datasets	Average Number of datasets	Average Standard Deviation (e-/s/pixel)	Average S/N per pixel
F098M	285	149	0.0026	263
F105W	113	51	0.0040	206
F110W	104	51	0.0029	239
F125W	374	179	0.0016	359
F140W	108	64	0.0046	137
F160W	934	483	0.0008	495
ALL/Grey	1894	976	0.0005	736

Table 3: WFC3 IR SD-flat properties.

Color Dependence

While, as Table 2 shows, we were not able to produce a high signal-to-noise SD-flat for all of the WFC3 filters, we can constrain the amount of wavelength dependence in the SD-flats using the lower signal to noise SD-flats.

Examining the fractional pixel to pixel difference between the F125W and the F160W filters, our two highest signal to noise monochromatic SD-flats, we compute that on average these two SD-flats differ by about +/-0.5%. In Figure 2 we show the fractional difference maps between these two SD-flats (Top panels, the panel on the right is further boxed averaged by 10x10 to be clearer. The bottom panel in Figure 2 shows the histogram of the fractional differences between the F125W and F160W SD-flats). Most of the field of view varies by less than +/-0.5%. The "wagon wheel" at the bottom right of the field is the most color dependent region of the detector, with a variation that is on the order of about 2%. We observe the same amount of variability between any pair of SD-flats taken from Table 1 and we see no clear indication of any color dependence in our correction to the CASTLE LP-flats. On this basis, we opted to build a single correction to the CASTLE LP-flats, a grey sky SD-flat constructed from all available data (nearly 2000 datasets). Figure 3 shows this SD-flat. It's properties are listed in Table 3 (ALL). Note that this does not mean that the WFC3 IR flats are not color dependent. The ground based CASTLE LP-flats do exhibit a significant amount of wavelength dependence. However, the in-orbit correction to the ground based flat-fields, likely caused by a difference in the illumination of the detector between CASTLE and HST, appears to be mostly wavelength independent.

Time Variation

We examined the temporal variation of the SD-flats in all filters by generating SD-flats using distinct time periods and sub samples of available data. For these tests, we used both monthly SD-flats as well as a set of SD-flats covering a period of 6 months each. Examining the fractional variation between these, we could not detect any temporal variation in the structure of the SD-flats, in any filter, and we estimate that any temporal variation to the SD-flats between 2009 and 2010 must be smaller than 1%.

Smoothing

As shown in Figures 3 and 4, the grey SD-flat provides some large scale correction to the ground based CASTLE LP-flats. This correction, as we have shown, is reasonably color independent and most likely caused by small differences between the CASTLE ground based testing bed and the actual HST optical train. Some of the pixel to pixel variation, i.e. very small scale length, that is present in the grey SD-flat is however the result of limited signal to noise. Our aim is to correct for the larger scale flat-field variations that were introduced by CASTLE. We opted to slightly de-noise the grey SD-flat using Fourier filtering and a filter with σ=10 pixels before combining the SD-flat with each individual, filter dependent CASTLE LP-flat. This filtering step was done to remove outlier pixels only. It has little or no impact otherwise.

Testing the SD-flats

HUDF Example

We provide a simple, real life example of the benefits of using the combined CASTLE LP-flat and grey SD-flat. We applied Multidrizzle to 32 F160W datasets from proposal 11563. The left panel of Figure 5 shows these data combined after being flatfielded using only the CASTLE LP-flats and the right panel shows the same data combined after being flat-fielded using the combination of CASTLE LP-flat and F160W SD-flat. The large slightly depressed cross pattern, offset towards the bottom-right side of the field of view, and visible in Figures 3 and 4, results in a significant amount of visual structure in the background of the left panel. The same structure is absent (<1%) in the combined data shown in the right panel.

Omega Cen Photometry

As a more quantitative test of the benefits of the SD-flats, we performed aperture photometry of Omega Cen data with and without the use of the grey SD-flat. We have identified stars that were observed in different regions of the detector (where the SD-flat correction differ the most) and compared their measured magnitudes with and without applying the SD-flat. We measured the brightness of stars in the Region 1 ([25:250,800:1000]) and Region 2 ([500:800,250:550]) on the detector (See Figure 3). We plot the difference in the measured magnitudes for 387 objects in Figures 6a and 6b with and without the use of the grey SD-flat. Figures 6a and 6b demonstrate

that the photometric accuracy is increased: for many objects, and the observed differences of about 3% are reduced to 1%. Overall, the uncertainty in photometry due to the flat-field correction has an estimated rms of <0.7% over the whole detector, with a maximum peak-to-peak range of -2.0/+1.9%. The uncertainties are typically larger at the edges of the detector. For the central part of the detector [129:896,129:896] (i.e., excluding a region around the detector edges with thickness 1/8 of the detector size), the uncertainty has an rms of <0.5%, with peak to peak range of -1.5/+1.6%. For the edge region of the detector only (i.e., a frame with 128 pixel thickness), which includes the "wagon-wheel" feature, the rms is <0.8%, with peak-to-peak -2.0/+1.9%. The introduced uncertainty in photometry due to the flat-fielding therefore typically has an rms of less than 0.01 mag, at the edges of the detector the errors may reach 0.02 mag. The uncertainty in the wavelength dependence of the sky image correction should also be less than 0.01 mag.

Self Consistency Test

As a final test to further estimate the amount of residual errors when applying the new flat-fields, we generated a new grey SD-flat using data that were already flat-fielded using the new CASTLE LP-flats+grey SD-flat. The resulting SD-flat is flat, as expected, with only a small amount of pixel to pixel variation. We measured the amount of residual information in this combined flat-fielded image by computing the running median and standard deviation over the whole the detector using a 10x10 pixels box. We did not include pixels with DQ bits of 4 (Bad detector pixel), 16 (Hot pixel), 32 (CTE tail), and 128 (Bad pixel in bias) when computing these statistics to avoid the effect of know bad pixels). The median of the combined flat-fielded data is very close to unity with a value of 1.0005 and varies little, with σ=0.001 e-/s. This is a good indication that the grey SD-flat, when combined to the individual CASTLE LP-flats, properly flattens WFC3 IR data. Within a 10x10 box, we measure that the pixel to pixel variation, computed all over the detector, is stable with a value of 0.0065 e-/s. This corresponds to a 0.5% residual noise level. This is taken to be a good indication that the CASTLE LP-flats+grey SD-flat provide an accurate flat-fielding of the data. Conservatively, new flat-fields that combine the ground based CASTLE data with the grey SD-flat provide flat-fielding accurate to better than 1%.

Comparison with photometry derived L-flats (stellar derived flats)

Photometry of the star cluster Omega Centauri was independently used to test the accuracy of the CASTLE LP-flats. Comparing our grey SF-flat to these stellar L-flats further allowed the photometric method to be vetted before it is used with the UVIS channel, for which generating SD-flats is not possible (due to the lower background levels at these shorter wavelengths). Calibration observations were obtained in program 11928 during Cycle 17 at 3 unique epochs (and roll angles): Dec 2009, Mar 2010, Aug 2010 (130 degrees, 194 degrees, 358 degrees). For each epoch, a 9 point dither pattern with a step size ~one quarter of the field of view was used to place stars on different regions of the detector with the goal of measuring low-frequency variations in detector response (L-flats), due to differences in the ground-based and inflight optical path.

The images were aligned and corrected for geometric distortion using MultiDrizzle, and aperture photometry was computed using daophot for each of the 27 images. A variety of apertures (r=1.5 to r=5) and sky annuli were tested to optimize the signal-to-noise of the IR photometry, while minimizing contamination from neighbors (at large apertures) or spatial variations in encircled energy (at small apertures). The best solutions were obtained using an aperture of 2 pixels in radius, with a sky annulus from 6-10 pixels. Approximately 1500 stars per image (with a signal-to-noise greater than ~25) were selected for analysis. While a signal-to-noise of 100 is ideal for constraining the L-flat at the 1% level, the IR exposures were not deep enough to achieve this goal while at the same time providing an adequate numbers of measurements over the detector.

The low-frequency residuals in inflight sensitivity were derived using the same software and methodology for the ACS L-flat corrections (see ACS ISRs 2003-10 and 2002-08). The residuals are shown for F110W and F160W in the top panel of Figure 7 with a stretch of +/- 2%, where black indicates that the photometry derived from images calibrated with the CASTLE LP-flats will be too faint. The detector was divided into a 16x16 grid, and a unique solutions was computed for each grid point. The bottom panel is a map of the number of measurements per grid point, where the average is ~150 and the minimum is ~5 measurements at the edges of the detector. While a 32x32 grid would give finer spatial sampling, the resulting solutions are much noisier do to 'holes' in the grid solution with very few measurements. The required L-flat corrections are ~1.1% rms, with a peak-to-peak of +/-2.5% in both filters. These L-flat solutions confirm the results from the sky flat both in the pattern of the spatial variation as well as it's amplitude, as well as the lack of a significant dependence on wavelength.

L-flats were also obtained from observations of the moonlit Earth lim in program 11917 (McCullough, private communication). Both the shape and amplitude of the low-frequency structures are strikingly similar to our SD-flats as well as those derived from dithered stellar observations of Omega-Cen. In Figure 8, the residuals are shown with a +/- 2% stretch the grey SD-flat (top left), the F105W Earth flat (top right), and the F160W stellar L-flat (bottom right). For comparison with the stellar L-flat, the grey SD-flat has been colorized and resampled to a 16x16 grid (bottom left). The small number of high signal-to-noise stars in the stellar L-flats, produces solutions which are significantly noisier than those obtained from other methods. As a result, these solutions were not used to correct the CASTLE LP-flats, but only to validate their results. The excellent agreement between these three independent tests verifies that the new SD-flat derived from stacks of sky images are appropriate both for observations sparse fields and for targets which uniformly illuminate the detector.

Data Release

New flat-fields that combine the CASTLE LP-flats and our new grey SD-flat were ingested into CDBS on December 7th, 2010. The file name and corresponding filter of these new flat-fields are listed in Table 3. The flats were created by dividing the existing CASTLE LP-flats by our grey SD-flat and normalizing them in the region [101:900:101:900].

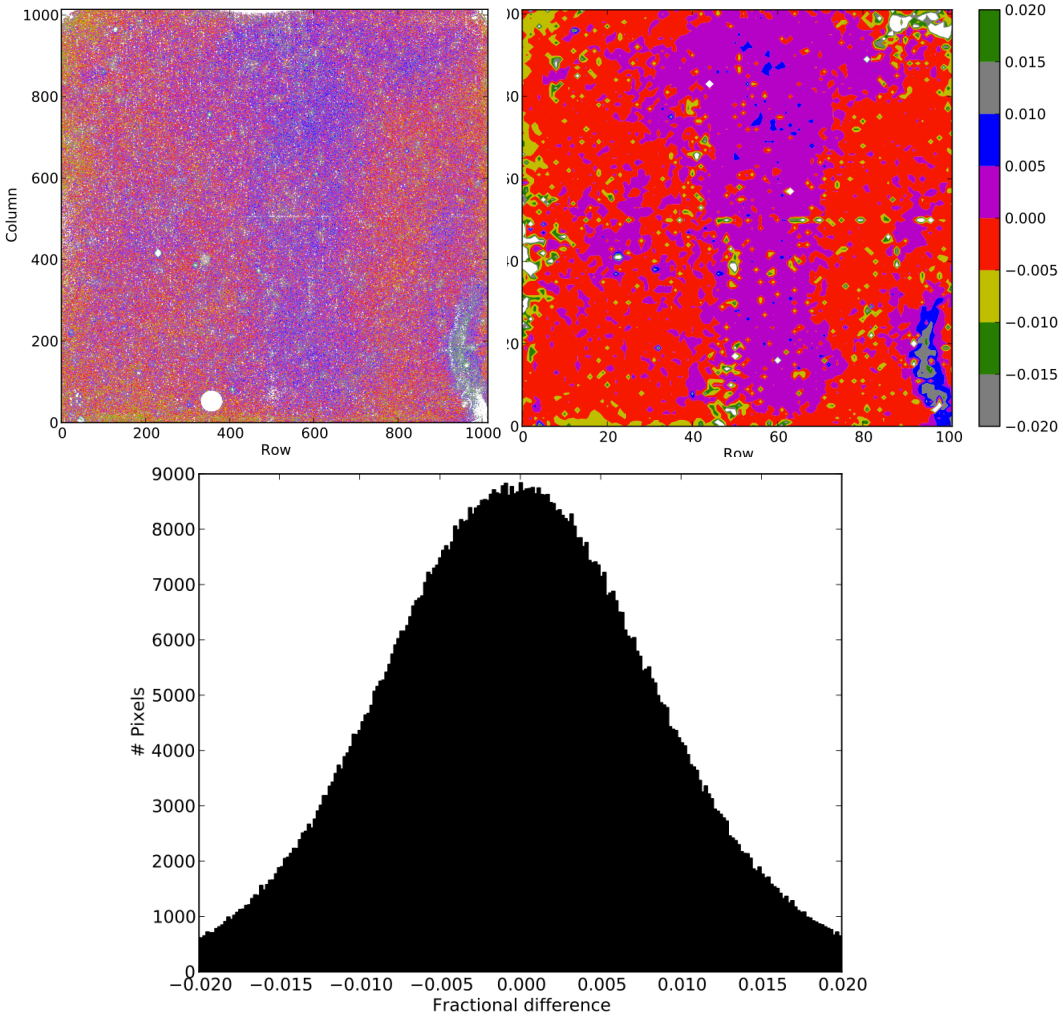

Figure 2: The fractional difference between the F125W and F160W SD-flats (top left panel). The top right panel is identical but box averaged (10x10 pixels) for clarity. The bottom histogram shows the distribution of fractional differences, which is closely approximated by a Gaussian centered on zero with full width of half maximum of about 0.01, or 1%.

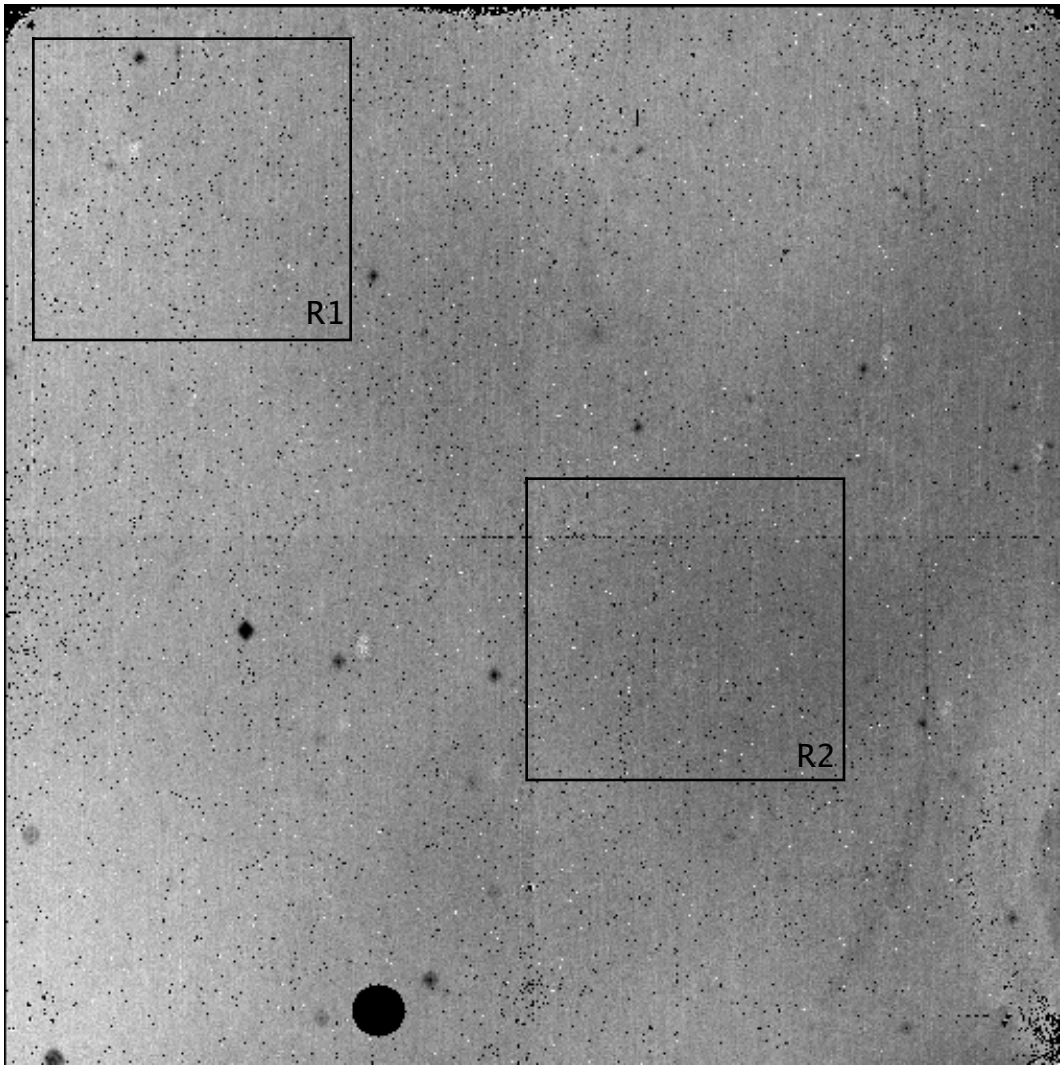

Figure 3: The grey IR SD-flat. The properties of this SD-flat are listed in Table 2. Blobs can be identified from this image (see ISR WFC3 2010-06). The wide cross like pattern is an illumination pattern that is present in the CASTLE LP-flat and that this SD-flat partially corrects. The same pattern can be seen in Figure 5 when combining CASTLE LP-flatfielded HUDF data. The black pixels are bad pixels. The large circular feature at the bottom of the detector is the WFC3 'Death Star', a known region of dead pixels. Both images are shown using the exact same stretch. The R1 and R2 regions described in the Testing section and Figures 6a and 6b are shown.

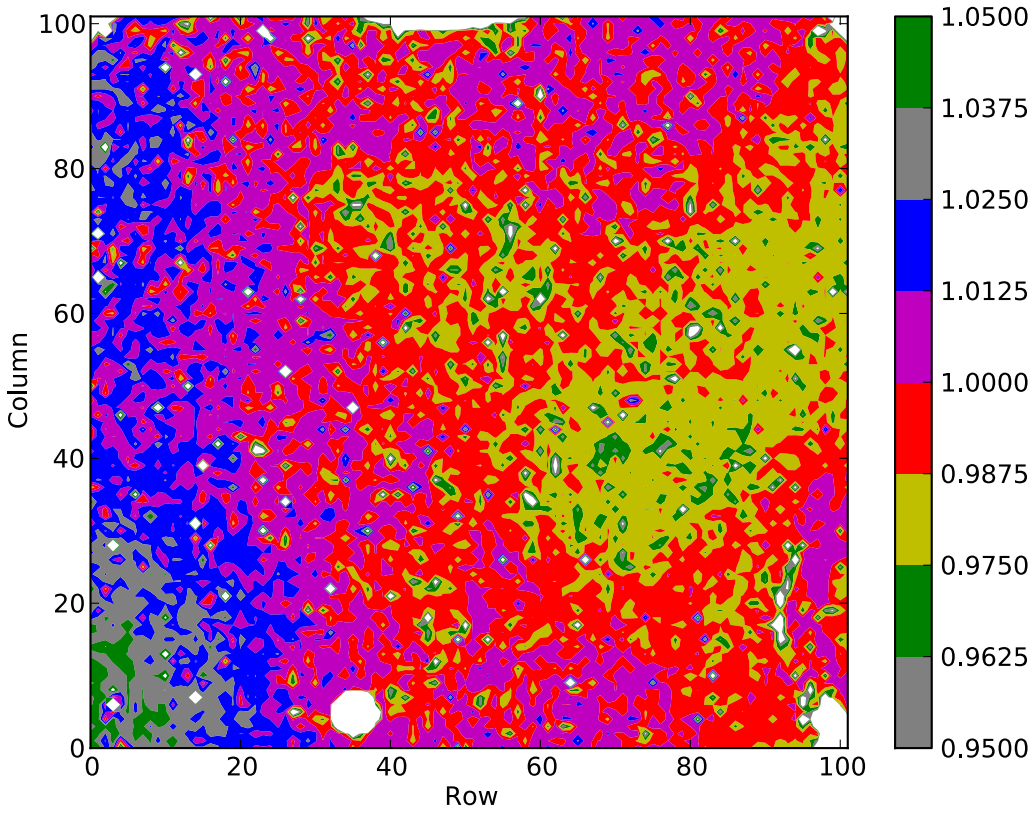

Figure 4: The SD-flat, slightly block averaged (10x10 pixels) and color coded to show the correction it provides to the CASTLE LP-flats. The color bar on the right shows the amount of correction. The SD-flat varies by about +/-2% over the most of the field of view.

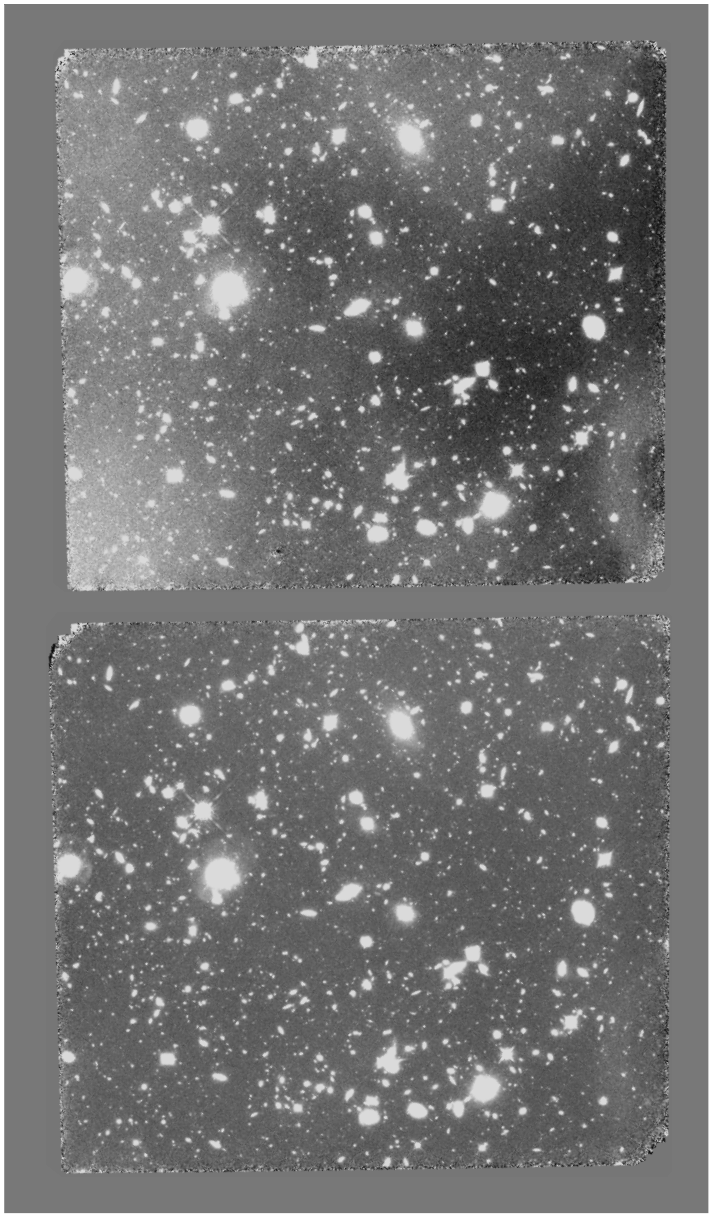

Figure 5: Results when applying multidrizzle to F160W observations of the HUDF. On the left panel we show the results when using CASTLE LP-flatfielded data. On the left panel we show the result obtained when using the new WFC3 F160W Sky flat (The combination of the grey SD-flat shown in Figure 3 and the existing CASTLE LP-flat). As expected, most of the background structure visible on the panel on the left is absent from the panel on the right.

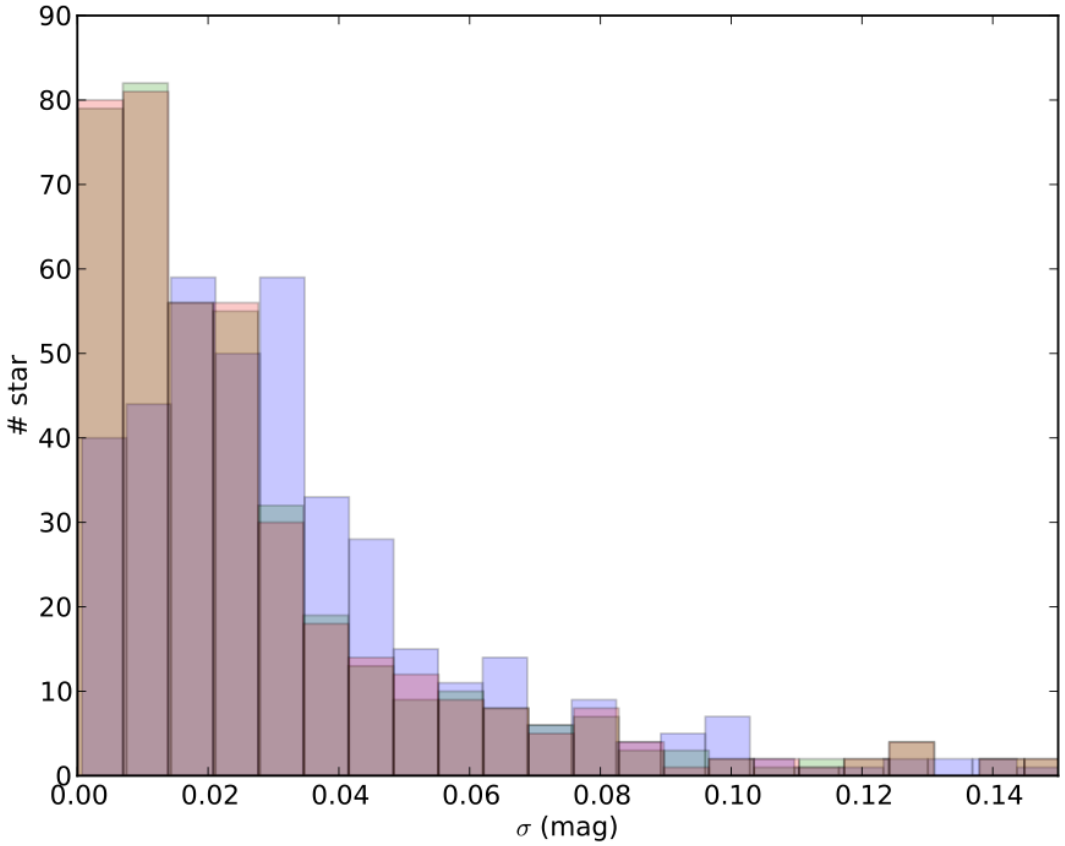

Figure 6a: Effect of flat-fielding on Omega Cen photometry. This shows the difference in measured magnitudes for stars that were in Region 1 and Region 2 on the detector. The blue semi-transparent histogram shows the observed magnitude difference when only the CASTLE LP-flat is applied. The green and red semi-transparent histograms show the same measurements when the F160W and grey SD-flats are used. These two histograms actually overlap almost completely and combine to make what looks like a beige histogram. The are of overlap between the green, red and blue distribution appears in purple. As this panel shows, typical photometric errors using the CASTLE LP-flat peaks around 0.03 (3%) magnitude (peak of the blue histogram). SD-flat corrected data result in photometric errors that peak at the lower value of 0.01 magnitude (1%) (peak of the non blue histograms).

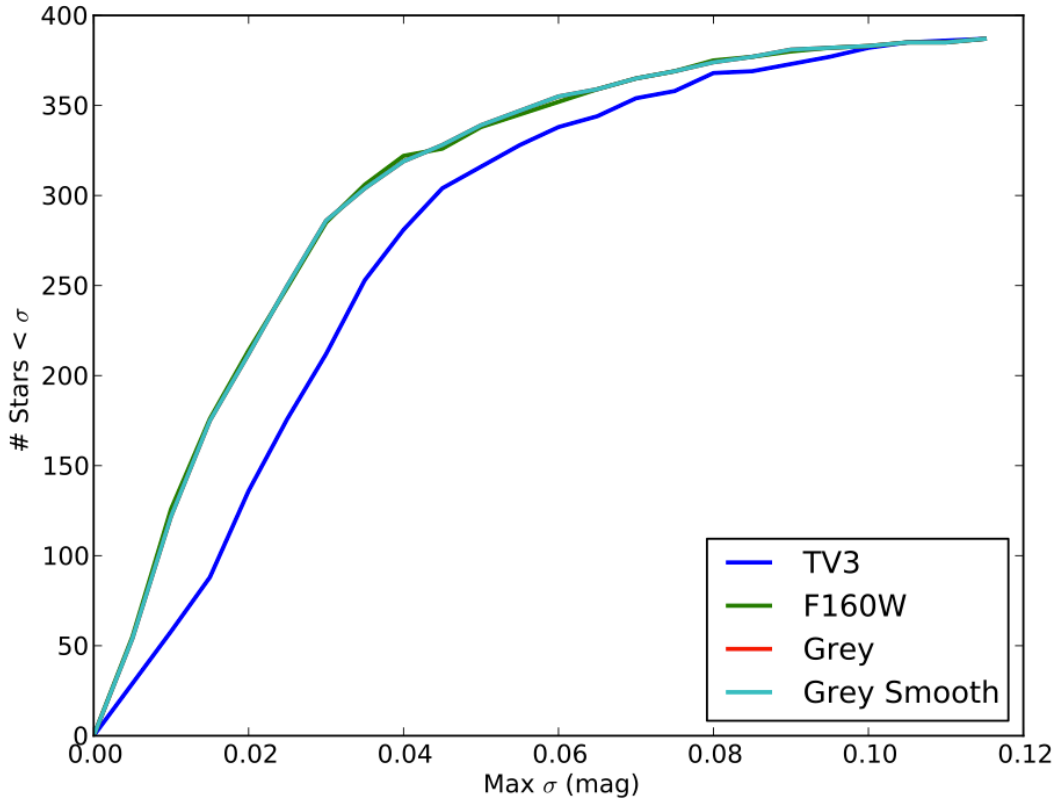

Figure 6b: This figure shows the same information as Figure 6a in cumulative form: the number of stars for which photometry differed by more than a given value. In blue (lower curve), we show the CASTLE LP-flatfielded data. In green, the result of using the F160W SD-flat. Finally, the red and cyan show the results of using the grey and de-noise/smoothed grey SD-flats. The red and cyan curves are not visible as they completely match the green curve (upper curves). This shows that the F160W, Grey and Smoothed Grey flat-fields affect photometry identically. This demonstrates the color independence of the correction to the CASTLE LP-flats.

Filter	Filename
F098M	uc72113ni_pfl.fits
F105W	uc72113oi_pfl.fits
F110W	uc72113pi_pfl.fits
F125W	uc72113qi_pfl.fits
F126N	uc72113ri_pfl.fits
F127M	uc72113si_pfl.fits
F130N	uc721140i_pfl.fits
F132N	uc721141i_pfl.fits
F139M	uc721142i_pfl.fits
F140W	uc721143i_pfl.fits
F153M	uc721144i_pfl.fits

Table 3: New WFC3 IR flat-fields, generated by combining the existing CASTLE LP-flats (ISR 2008-28) and the grey SD-flat described in this ISR.

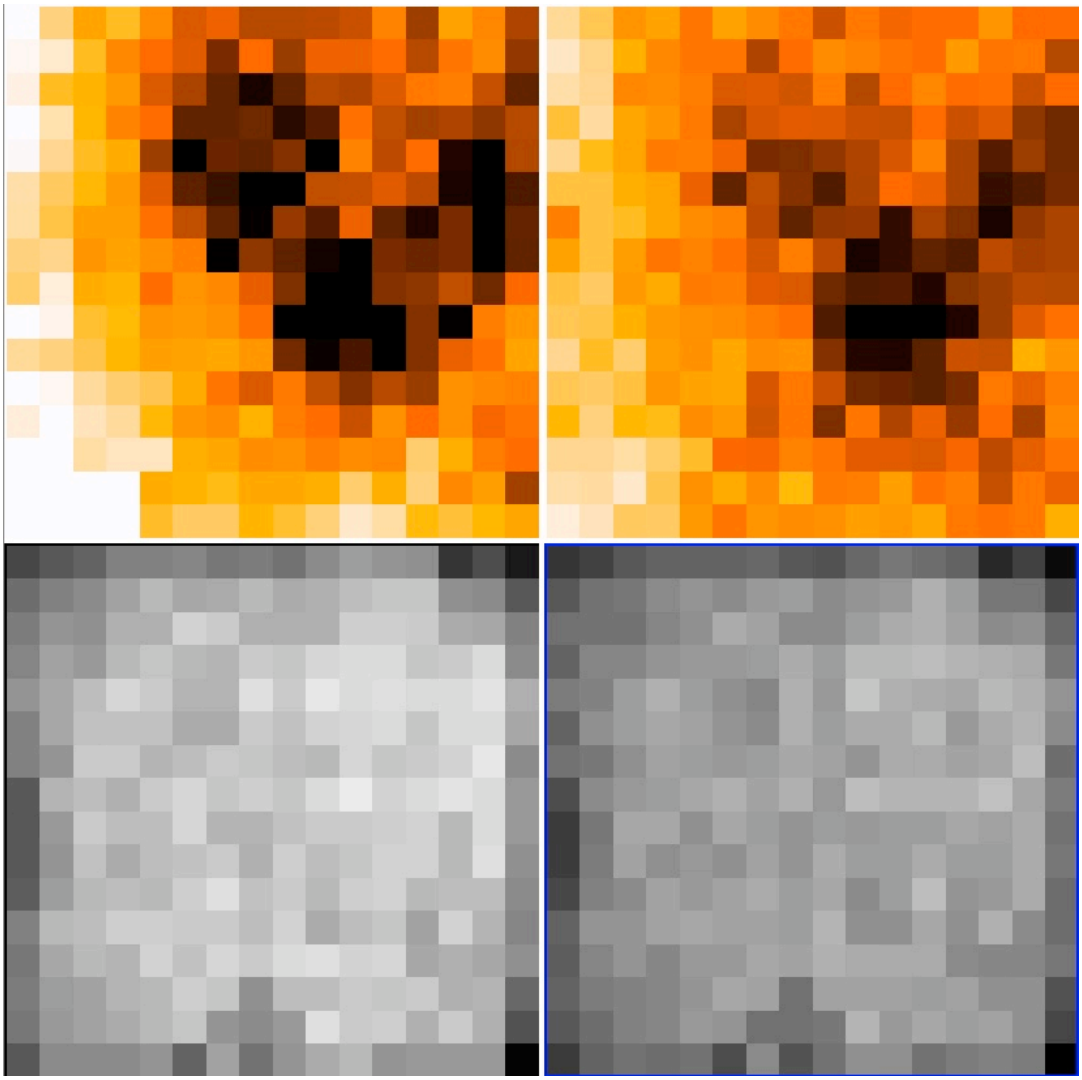

Figure 7: Top Panel: F110W (left) and F160W (right) Lflat solutions. Bottom Panel: F110W and F160W weight map (number of measurements per grid point).

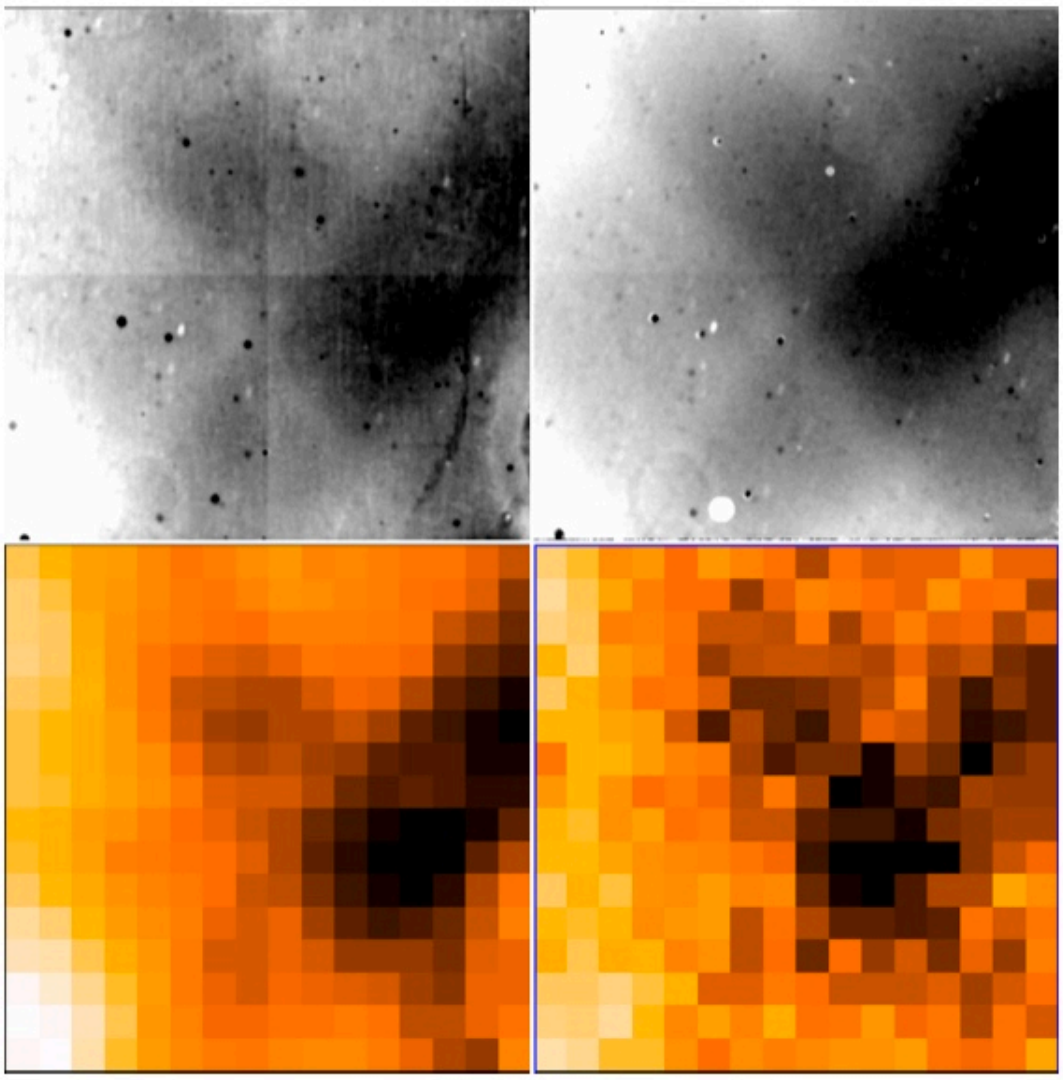

Figure 8: Top Left = Grey SD-flat. Top right = Earth flat (F105W). Bottom Left= 16x16 binned sky flat (colorized). Bottom right = Stellar L-flat from aperture photometry of Omega Cen (F160W).