7.2 WFC3/IR: Optimizing Image Sampling for a Single Visit

Introduction

This example was written to help users better understand the subtleties in improving image sampling for dithered data. Four images of a spiral galaxy were acquired using WFC3/IR, following the `WFC3-IR-DITHERBOX-MIN` dither pattern that was designed to provide optimal sampling of the PSF.

Summary of Steps

1. Description of the Data
2. Run **astrodrizzle** several times using different settings for the ***final_pixfrac*** and ***final_scale*** parameters
3. Compare and evaluate results of using different ***final_pixfrac*** and ***final_scale*** values

7.2.1 Description of the Data

Four WFC3/IR images[4] of the spiral galaxy NGC 3370 (Program 11570), taken in the F160W filter, were acquired in a single visit and at the same telescope orientation. Observations were obtained using the default WFC3/IR dither pattern, `WFC3-IR-DITHERBOX-MIN,` with relative pixel coordinates (0, 0), (4.2, 1.4), (2.6, 3.8), (-1.6, 2.4), which is designed to provide optimal PSF sampling.

Calibrated data products from the Archive are:

- An association table, with suffix `asn.fits`
- Flat field-calibrated images, with suffix `flt.fits`
- Drizzled image product, with suffix `drz.fits` that was created by running AstroDrizzle in the pipeline with a default set of parameters.

The pipeline `drz.fits` image may be saved to a separate directory for later comparison with the drizzled products from this example. In general, drizzled data from the Archive should be regarded as "quick look" data products, used to make an initial evaluation of the observations.

4. Data for this example can be retrieved from the HST Archive by searching for Dataset `ib1f19010`.

Table 7.3: Summary of Images in this Example

Image Name	Association ID	Proposal ID	Visit & Line Number	POS TARG)x,y in arcseconds)	PA_V3 Orientation (degrees)	Observation Date	Exposure Time (sec.)
ib1f19l6q_flt.fits	IB1F19010	11570	19.001	0.0000,0.0000	320.9999	2010-04-04	502.9365
ib1f19l7q_flt.fits	IB1F19010	11570	19.001	0.5423, 0.1818	320.9999	2010-04-04	502.9365
ib1f19l9q_flt.fits	IB1F19010	11570	19.001	0.3389,0.4848	320.9999	2010-04-04	502.9365
ib1f19laq_flt.fits	IB1F19010	11570	19.001	-0.2034,0.3030	320.9999	2010-04-04	502.9365

Since the data were obtained in a single visit as part of a subpixel dither box pattern, the WCS of the individual frames are usually aligned to 0.1 pixels. This example does not describe the use of TweakReg to verify (and/or improve) image alignments, but users are strongly encouraged to do so because even the smallest misalignment can compromise the photometric integrity of the final drizzled products.

For the IR detector[5], calibrated data products (`flt.fits`) consist of five extensions:

- science image (SCI)
- error array (ERR)
- data quality array (DQ)
- number of samples array (SAMP)
- integration time array (TIME)

A WFC3/IR FITS file will therefore contain the primary header unit and five extensions, which together form a single IR exposure. To see the contents of the IR file structure, the user can use the IRAF task **catfits**, shown below in a PyRAF session.

```
--> catfits ib1f19l6q_flt.fits
EXT#  FITSNAME       FILENAME  EXTVER DIMENS     BITPIX
0     ib1f19l6q_flt  ib1f19l6q_flt.fits          16
1     IMAGE          SCI       1      1014x1014  -32
2     IMAGE          ERR       1      1014x1014  -32
3     IMAGE          DQ        1      1014x1014  16
4     IMAGE          SAMP      1      1014x1014  16
5     IMAGE          TIME      1      1014x1014  -32
```

5. For details on the IR channel file structure, see Section 2.2.2 of the *WFC3 Data Handbook* at http://www.stsci.edu/hst/wfc3/documents/handbooks/currentDHB/wfc3_Ch23.html#96833

The SCI, DQ, and TIME extensions are shown in Figures 7.9, 7.10, and 7.11 for the first flt.fits image in the association. They were displayed using these IRAF commands:

```
--> display ib1f19l6q_flt.fits[sci,1]  1 zs- zr- z1=0.50 z2=100 ztr=log fill+
--> display ib1f19l6q_flt.fits[dq,1]    2 zs+ zr+ fill+
--> display ib1f19l6q_flt.fits[time,1] 3 zs+ zr+ fill+
```

The TIME extension is useful for identifying pixels which were saturated in one or more samples (such as the core of the galaxy in this example, which saturated after 8 of 12 total samples), or cosmic rays which were flagged in "up-the-ramp" fitting, usually in a single sample. (For more information, see Section 3.4.3 of the *WFC3 Data Handbook*.) Note that cosmic rays flags[6] are actually flagged in the IMA[7] files with a bit value of 8192, but their effect can be seen in the "reduced" exposure time in the 5th extension of the flt.fits files.

Figure 7.9: Science Portion of the Calibrated WFC3 Image

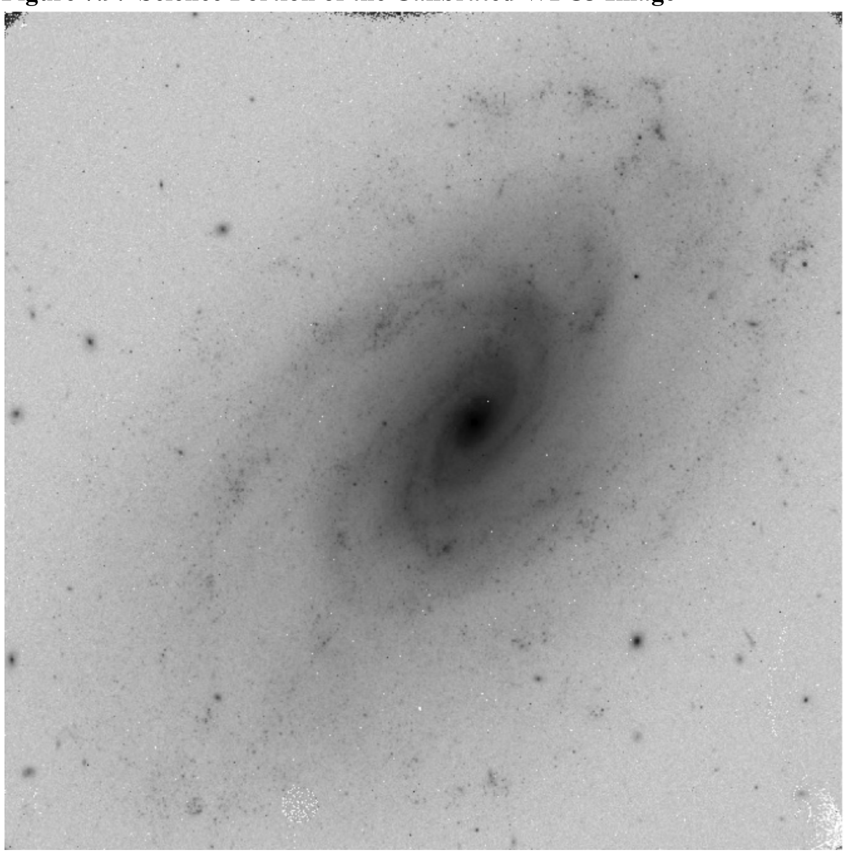

6. The specific DQ flag values are unique for each detector and are defined in the Instrument Data Handbooks. For a table of WFC3/IR DQ flags, see Table 2.5 in the *WFC3 Data Handbook* at http://www.stsci.edu/hst/wfc3/documents/handbooks/currentDHB/wfc3_Ch23.html#98193

7. For more information about WFC3 data products, see Section 2.1.1 in the *WFC3 Data Handbook* at http://www.stsci.edu/hst/wfc3/documents/handbooks/currentDHB/wfc3_Ch22.html#96161

Figure 7.10: Data Quality Portion of the Calibrated WFC3 Image

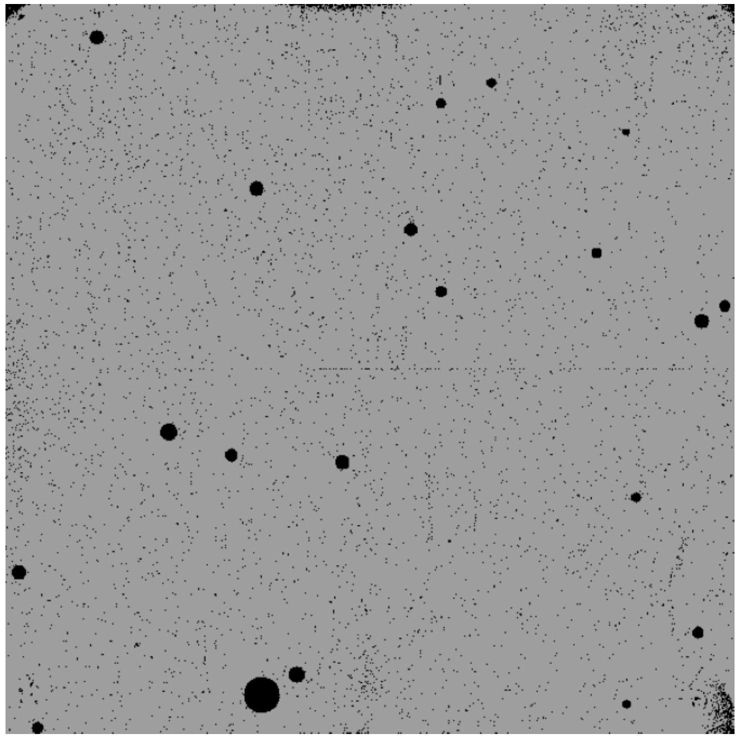

Figure 7.11: Exposure Time (TIME) Extension of the Calibrated WFC3 Image

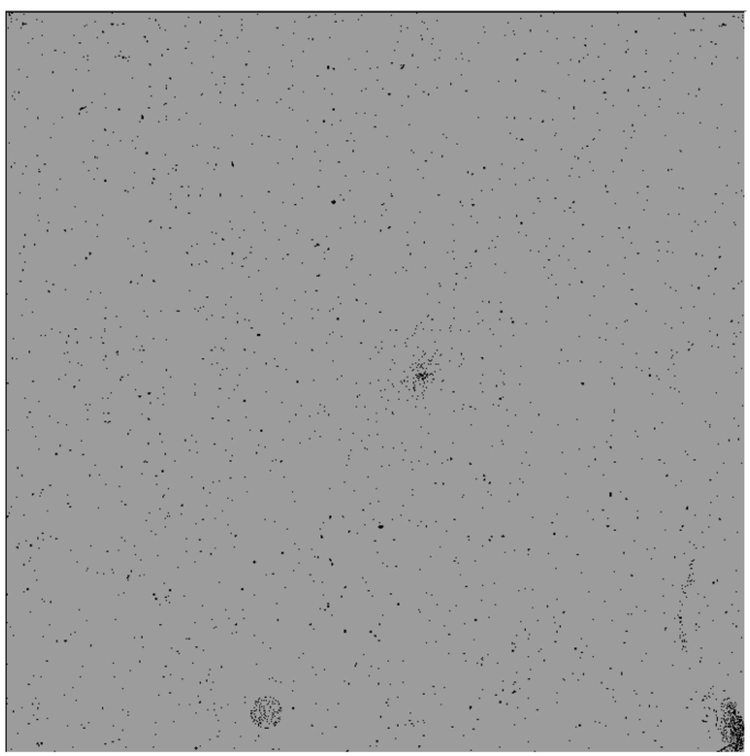

Black pixels indicate a lower total exposure time. The galaxy core, just off the center of the frame, was saturated after 8 of the 12 total samples.

7.2.2 Recommendations on Selecting the Optimal "Scale" and "Pixfrac" Parameter values

To optimize the parameters for drizzle combination, users are encouraged to experiment with various combinations of the parameters ***final_scale*** (size, in arcseconds, of the output pixels) and ***final_pixfrac*** (the fractional linear size of the input pixel "drop" into the output image frame).

The recommended method is to first select the ***final_scale*** value, then the ***final_pixfrac*** value. While experimenting with ***final_scale***, the ***final_pixfrac*** should be fixed at 1.0. Ideally, the scale is chosen to sample the PSF FWHM by about ~2.0 to 2.5 pixels, if allowed by the data. Non-integral (subpixel) dithers allow the recovery of some information lost to undersampling by pixels that are large compared to the point spread function. The FWHM of the IR point spread function is approximately 1.0 pixel, so subpixel dithering allows the user the ability to recover spatial resolution. (Ideally one would like a minimum two samples per FWHM for the full recovery of the image resolution.)

While reducing the ***final_scale*** from the default value, the PSF will begin to degrade and resemble the dither pattern (a "cross-shaped" PSF, for example, for a four-point dither). This is illustrated in Figure 7.12, where ***final_scale*** has been decreased from 0.1283 arcseconds/pixel to 0.0642 arcseconds/pixel, then to 0.032 arcseconds/pixels, while maintaining ***final_pixfrac*** at 1.0. In general, the ***final_scale*** value should never be less than half the native plate scale.

The task **imexamine** was used to plot both the radial profile (top panel) and a contour plot (bottom panel) of a bright star in the drizzled image, where the FWHM is 1.6 pixels, 2.8 pixels, and 6.2 pixels, respectively. With a well-sampled 4-point subpixel dither, the best ***final_scale*** is approximately half the native scale. Often, a "convenient" number is chosen, for example, 0.065 arcseconds/pixel.

Alternately, when UVIS and IR images are obtained with four-point dithering, the former scale could be set to 0.03333 arcseconds/pixel and the later set to 0.06666 arcseconds/pixel, a factor of two difference. For WFC3/UVIS (and ACS/WFC) an output scale is 0.03333 arcseconds/pixel gives good subsampling of the PSF. It is not quite a factor of two smaller than the original pixel (which would essentially recover all of the fine scale information in the image) but tends to sample the PSF very well. For a dozen or more pointings well distributed over the image (not just a multiple repetition of a four-point dither) a finer output pixel scale could be used if high-resolution imaging is important. A scale of 0.03333 arcseconds/pixel has the virtue that three pixels is ~0.1 arcseconds, making it easy to look at the output image and know the size of an object. When WFC3/UVIS (or ACS/WFC) images are obtained with corresponding WFC3/IR observations, it may be "convenient" to select a final scale which is a factor of two larger, where three pixels is ~0.2 arcseconds.

Figure 7.12: Radial Profile and Contour Plot of a Bright Star in Images with Three Different *final_scale* Values

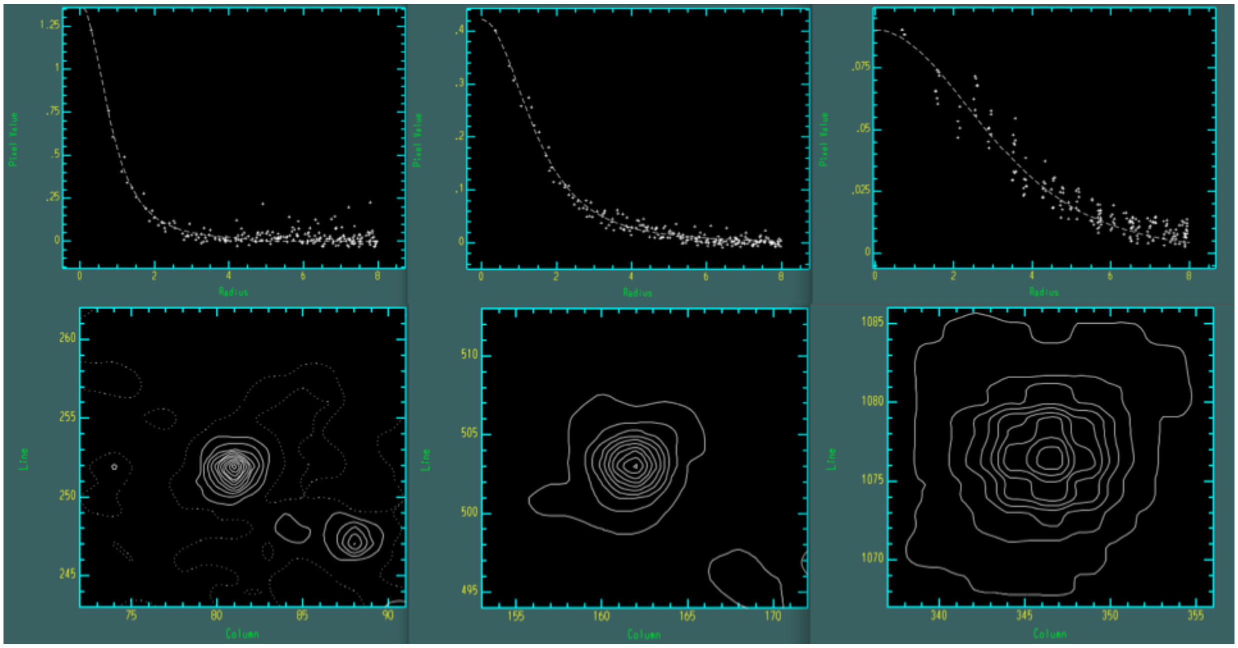

Radial profile and contour plot of the same star in drizzled frames obtained when *final_scale* is "shrunk" from 0.1283 arcseconds/pixel to 0.0642 arcseconds/pixel, to 0.0320 arcseconds/pixel. When the scale is set too small, the PSF shape begins to resemble the four-point dither pattern used in this observing program. The optimal scale value ultimately depends on the dataset, the number of dithers, and the amount of subpixel sampling. In this case, the middle panel is ideal.

The ***final_pixfrac*** value has to be small enough to avoid degrading the final drizzle-combined image, but large enough that when all images are "dropped" onto the final frame, coverage of the output frame is fairly uniform. In general, ***final_pixfrac*** should be slightly larger than the final output scale to allow some "spillover" to adjacent pixels. This will help avoid "holes" in the final product when a given pixel has been flagged as "bad" in several frames. As a rule of thumb, statistics performed on the drizzled weight image in the region of interest should yield an RMS value (standard deviation) that is less than 20% of the median (midpoint) value. This threshold is a balance between the benefits of improving the image resolution at the expense of increasing noise in the background.

7.2.3 Image Combination with AstroDrizzle

In default mode, AstroDrizzle performs each of its seven steps in the order outlined in Section 4.2. For IR images, however, steps three to six may be turned off since cosmic rays are flagged in **calwfc3** as part of the "up-the-ramp fitting." While it is omitted from this specific example, running these steps (using a different bit flag, like 8192, for "cosmic rays" found during **astrodrizzle** processing) may still be useful for flagging additional detector artifacts not present in the data quality arrays of the calibrated images. Note that it is very important to subtract the sky (step two) prior to drizzing the final image, or the science array will be compromised by increased noise. The size of the effect will depend on the variation in the sky between exposures. (An example of this effect is shown in Figure 7.16)

The commands shown below run a test grid of varying ***final_scale*** and ***final_pixfrac*** values to show how the images change at different settings.

When the parameter ***build=yes*** (a non-default value), the final AstroDrizzle output image for this example will be a single multi-extension FITS file named `f160w_drz.fits`, containing the science image in extension one, the weight image in extension two, and the context image in extension three. When ***build=no***, the science, weight, and context images are written to separate output files. Since the output file `f160w_drz.fits` will be overwritten with each successive run, this example renames the drizzled product with a unique name between each separate trial. The commands below use the command-line syntax; non-default parameter values are highlighted in bold.

```
--> import drizzlepac
--> from drizzlepac import astrodrizzle
--> unlearn astrodrizzle
--> astrodrizzle.AstroDrizzle('*flt.fits',output='f160w',build=yes,\
static=no,skysub=yes,driz_separate=no,median=no,blot=no,driz_cr=no,\
driz_combine=yes,final_wcs=yes,final_bits=576,final_scale=0.1283,\
final_pixfrac=1.0)
--> imrename f160w_drz.fits f160w_drz_test1.fits
```

Next, run the previous **astrodrizzle** commands, varying only the ***final_scale*** parameter.

```
--> astrodrizzle.AstroDrizzle('*flt.fits',output='f160w',build=yes,\
static=no,skysub=yes,driz_separate=no,median=no,blot=no,driz_cr=no,\
driz_combine=yes,final_wcs=yes,final_bits=576,final_scale=0.0898,final_pixfrac=1.0)
--> imrename f160w_drz.fits f160w_drz_test2.fits

--> astrodrizzle.AstroDrizzle('*flt.fits',output='f160w',build=yes,\
static=no,skysub=yes,driz_separate=no,median=no,blot=no,driz_cr=no,\
driz_combine=yes,final_wcs=yes,final_bits=576,final_scale=0.0642,final_pixfrac=1.0)
--> imrename f160w_drz.fits f160w_drz_test3.fits

--> astrodrizzle.AstroDrizzle('*flt.fits',output='f160w',build=yes,\
static=no,skysub=yes,driz_separate=no,median=no,blot=no,driz_cr=no,\
driz_combine=yes,final_wcs=yes,final_bits=576,final_scale=0.0513,final_pixfrac=1.0)
--> imrename f160w_drz.fits f160w_drz_test4.fits
```

Once the scale is chosen, the value of ***final_pixfrac*** may then be varied. Note that while ***final_scale*** is represented in arcseconds, ***final_pixfrac*** is represented as a fraction of the native pixel size.

```
--> astrodrizzle.AstroDrizzle('*flt.fits',output='f160w',build=yes,\
static=no,skysub=yes,driz_separate=no,median=no,blot=no,driz_cr=no,\
driz_combine=yes,final_wcs=yes,final_bits=576,final_scale=0.0642,final_pixfrac=0.9)
--> imrename f160w_drz.fits f160w_drz_test5.fits

--> astrodrizzle.AstroDrizzle('*flt.fits',output='f160w',build=yes,\
static=no,skysub=yes,driz_separate=no,median=no,blot=no,driz_cr=no,\
driz_combine=yes,final_wcs=yes,final_bits=576,final_scale=0.0642,final_pixfrac=0.8)
--> imrename f160w_drz.fits f160w_drz_test6.fits

--> astrodrizzle.AstroDrizzle('*flt.fits',output='f160w',build=yes,\
static=no,skysub=yes,driz_separate=no,median=no,blot=no,driz_cr=no,\
driz_combine=yes,final_wcs=yes,final_bits=576,final_scale=0.0642,final_pixfrac=0.7)
--> imrename f160w_drz.fits f160w_drz_test7.fits

--> astrodrizzle.AstroDrizzle('*flt.fits',output='f160w',build=yes,\
static=no,skysub=yes,driz_separate=no,median=no,blot=no,driz_cr=no,\
driz_combine=yes,final_wcs=yes,final_bits=576,final_scale=0.0642,final_pixfrac=0.6)
--> imrename f160w_drz.fits f160w_drz_test8.fits
```

The first extension of the drizzled product, `160w_drz.fits[1]` contains the science (`SCI`) image, a combination of the four dithered images which has been corrected for distortion. All pixels cover an equal area on the sky and have an equal photometric normalization across the field of view, giving an image that is photometrically and astrometrically accurate for both point and extended sources. The `SCI` portion of the drizzled product, shown in Figure 7.13, is in units of electrons/seconds. (Changing the ***final_units*** parameter from the default value ***cps*** (counts per second) to ***counts*** will produce a drizzled image in units of electrons.)

The second extension of the output image contains the weight (`WHT`) image. When ***final_wht_type*** is set to ***EXP***, the weight image can be considered an effective exposure time map of the science (`SCI`) image. In Figure 7.14, darker areas in the `WHT` extension image have lower weights. IR weight images represent several different types of information; when ***final_pixfrac***=***1.0***, the weight image will resemble the `TIME` extension of the `flt.fits` image, minus the pixels which were flagged in the flt.fits DQ array and not specifically set as "good" in the final_bits parameter.

Note that in this example, the ***final_bits*** value is ***576*** (it can also be written as ***512,64***) to tell **astrodrizzle** that `flt.fits DQ` flags of 512 (bad pixels in the flat field) and 64 (warm pixels) should be treated as "good" pixels. All other DQ flags in the `flt.fits` images, treated as "bad," are reflected in the `single_wht.fits` weight images. The smooth top-to-bottom gradient in the `WHT` image reflects the geometric distortion in the IR detector where detector pixels represent different areas on the sky. When ***final_pixfrac*** is shrunk to values smaller than ***1.0***, the RMS of the `WHT` image increases, as shown in the bottom panel of Figure 7.14. When ***final_pixfrac*** is too small relative to `final_scale`, there will be pixels with "holes" in the weight image where less than one pixel contributed to the value of the final flux in the drizzled science image. The majority of the variations in the bottom `WHT` image is due to the change in geometric distortion over the chip, where the input pixels cover significantly different areas on the sky.

Figure 7.13: Science Extension of the Final WFC3 Drizzled Product

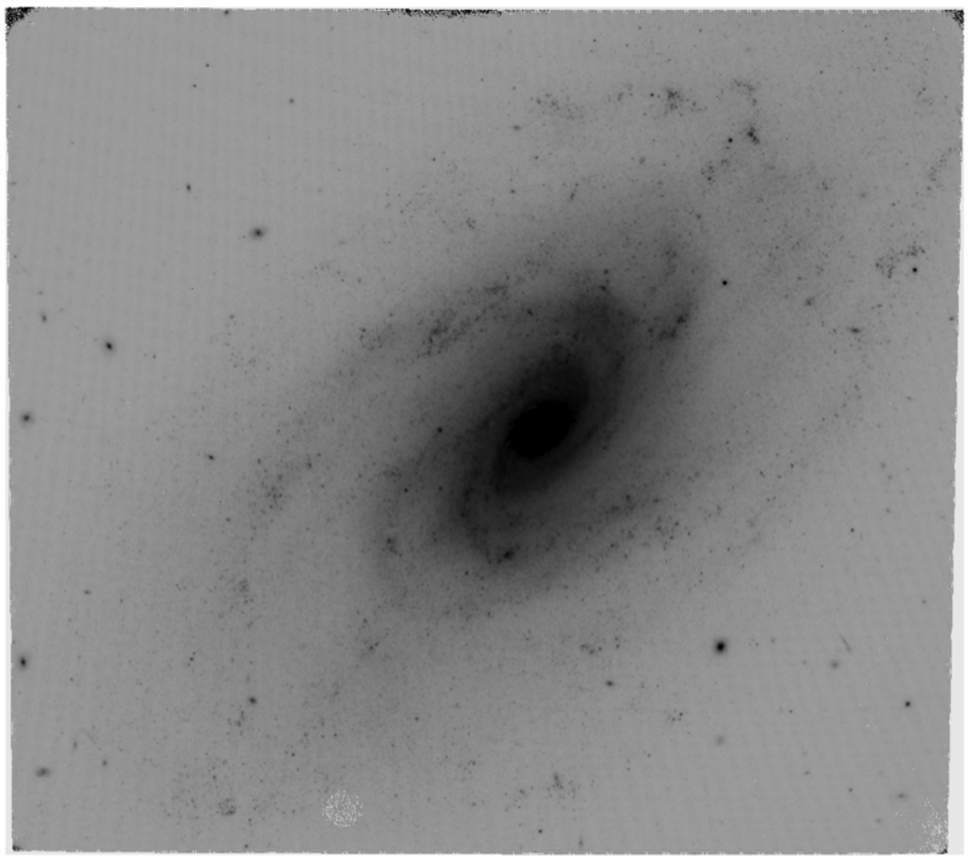

Figure 7.14: Weight Extension of the Final WFC3 Drizzled Product (Top: *final_pixfrac=1.0*; **Bottom:** *final_pixfrac=0.8*

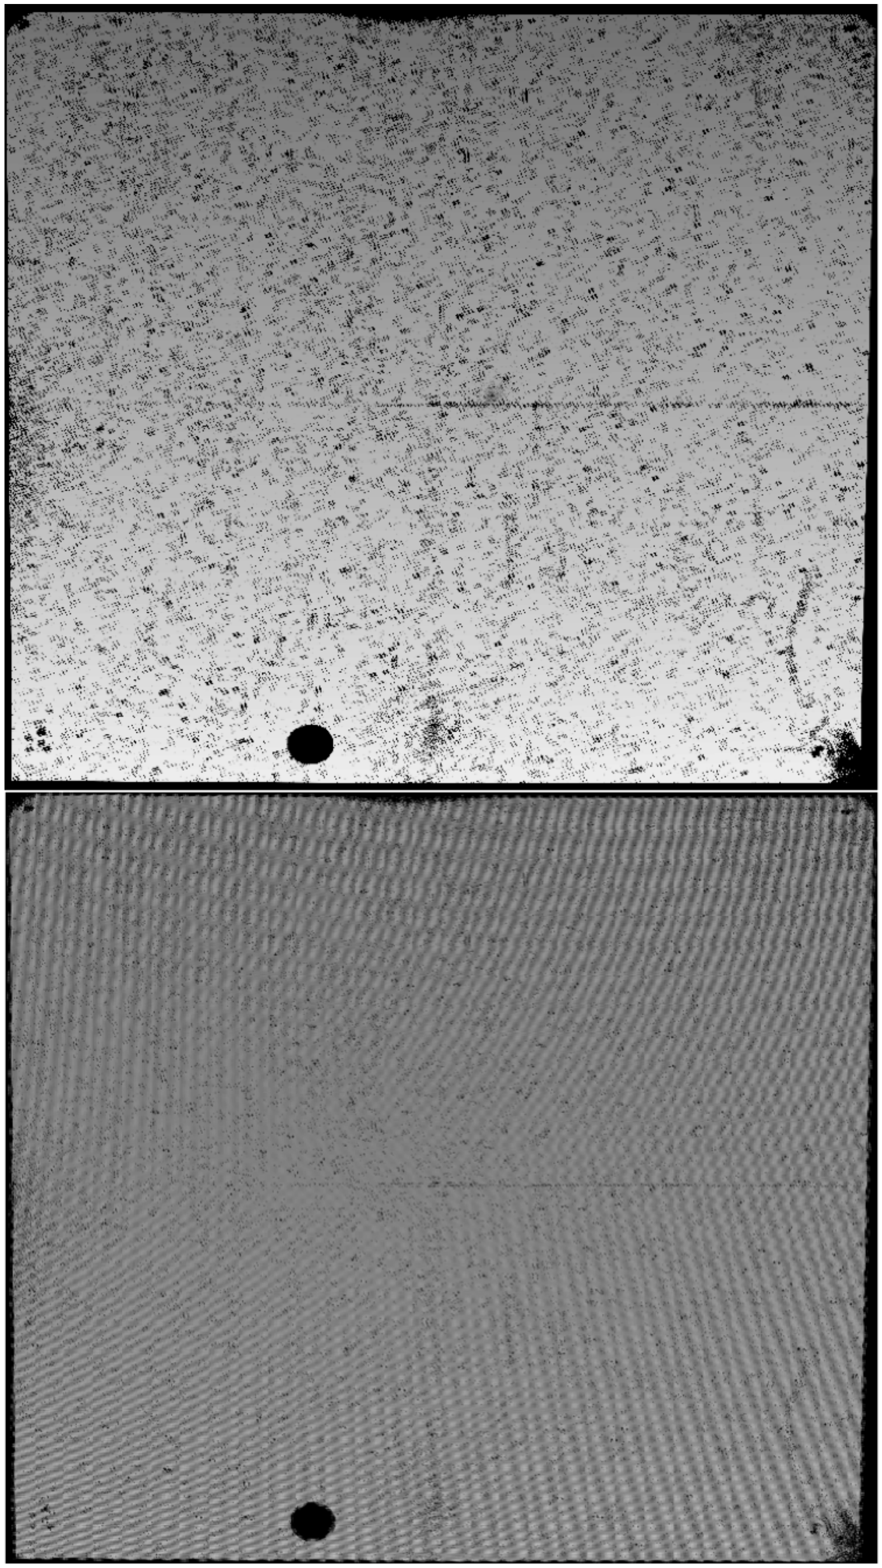

Statistics in the weight image (RMS/median) are reported in Table 7.4 for several trials, computed using the IRAF task **imexam** for a 200x200 pixel box in the center, and in the top left corner of each weight image. The PSF FWHM was measured using an isolated star at coordinate (430, 1746) in the trial image, where the ***final_scale*** setting used to create it is ***0.0642*** (in arcseconds/pixel). Note that the table gives the value of ***final_scale*** (shown in the table as “Scale”) in two different ways: as a fraction of the default plate scale and in arcseconds/pixel (the units used in **astrodrizzle**).

Table 7.4: Weight Image Statistics and PSF FWHM for Various Final Drizzle Scale/Pixfrac Combinations

Trial Number	Pixfrac (fraction)	Scale (fraction)	Scale (arcsec.)	RMS/Median (center)	RMS/Median (corner)	PSF FWHM (pixels)	PSF FWHM (arcsec.)
3	1.0	0.5	0.0642	0.061	0.066	2.97	0.191
5	0.9	0.5	0.0642	0.068	0.073	2.90	0.186
6	0.8	0.5	0.0642	0.076	0.076	2.85	0.183
7	0.7	0.5	0.0642	0.083	0.095	2.78	0.179
8	0.6	0.5	0.0642	0.090	0.109	2.70	0.173

Statistics of the weight image for both regions of the detector meet the general guideline of rms/median < 0.2 for all the trials. However, if one visually compares the science products in Figure 7.15, it becomes apparent that maintaining a larger ***final_pixfrac*** ensures overlap between pixels and less correlated noise in the science array (bottom panel). When ***final_pixfrac*** has been shrunk too much (top panel), a “beating pattern” can be seen in the sky. While this pattern may look alarming to the eye, it has only a very minor effect on the photometric integrity of the drizzled products.

Determining which is the best solution is a matter of judgment, depending on the preferred resolution and quality of image. If the target is primarily in the center of the frame, the ***final_scale*** and ***final_pixfrac*** selection may be more aggressive. If sources cover the entire field of view, however, a more conservative set of parameters may be preferable.

While trial number eight gives a narrower PSF FWHM, it does not do a good job at removing detector artifacts. Trial number 6 is shown in Figures 7.13 and 7.14, where the ***final_scale*** is equal to 0.5 times the default pixel scale and ***final_pixfrac*** is ***0.8***. The resulting image has a plate scale of 0.0642 arcseconds/pixel with the PSF FWHM at 0.183 arcseconds. Because the WFC3/IR detector pixels are significantly undersampled, optimizing the ***final_scale*** and ***final_pixfrac*** parameters will produce a dramatic improvement in resolution, as seen in Figure 7.17.

Figure 7.15: Comparison of *final_pixfrac=0.6* (top) and *0.8* (bottom)

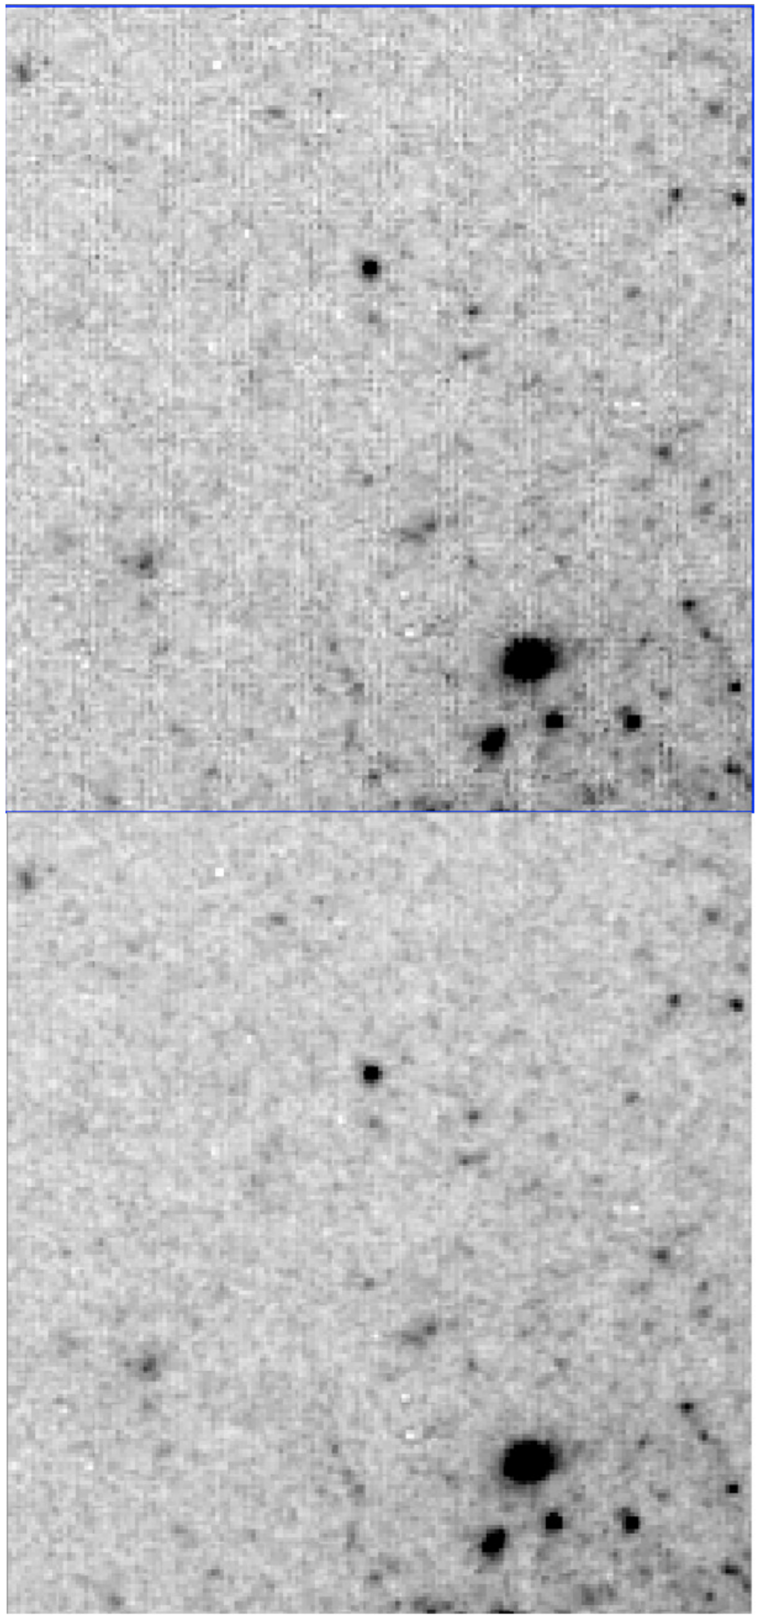

Figure 7.16: Sky Background in the Final Science Array With No Sky Subtraction (Top) and With Sky Subtraction (Bottom)

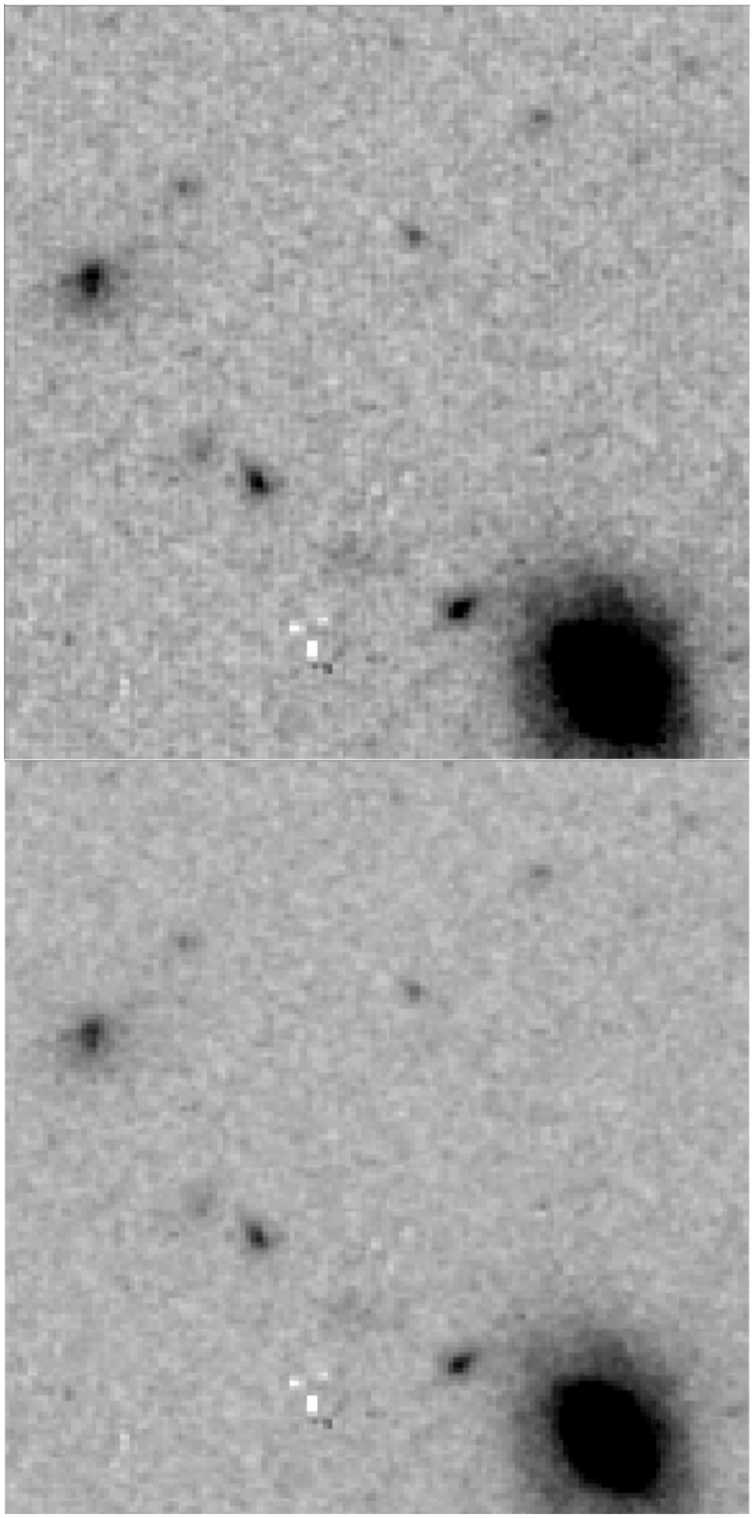

Note the additional noise in the top panel. The effect in this example is subtle, but will be more pronounced in images with larger sky variability between exposures.

Figure 7.17: Improvement in Resolution of the Pipeline Product (Top) Versus the Optimized Drizzled Product (Bottom) I