

CANDELS HST Imaging Data README

Anton M. Koekemoer and the CANDELS Team

Data Summary

Release Version	v1.0
CANDELS Field	COSMOS
Subregion	all
Epoch	cos 01 & 02
Field Tangent Point	R.A. = 150.1163213, Dec. = +2.200973097 (J2000)
Area	210 sq arcminutes
Program ID	12440
Epoch Observation Dates	2011-Dec-06 – 2012-Apr-16
Telescope	HST

Instrument	ACS/WFC		WFC3/UVIS	WFC3/IR	
Mosaic pixel Scale	0'03		0'03	0'06	
Filter	F606W	F814W	F350LP	F125W	F160W
PSF FWHM	0'08	0'09	0'08	0'12	0'18
Exposure Depth	3300s	6900s	434s	1900s	3200s
Zeropoint	26.493	25.947	26.94	26.2303	25.9463
Sensitivity (5σ pt. src.)	28.5	28.4	26.4	27.0	26.9

Flux Unit e/s

The CANDELS overview/data reference papers:

- [Grogin, N. A. et al., 2011 ApJS, 197, 35](#) (“CANDELS Survey and Observational Design”)
 - [Koekemoer, A. M. et al., 2011, ApJS, 197, 36](#) (“CANDELS HST Imaging Data Products”)
-

Contents

1	General Information	3
2	Data Products	4
2.1	Science Images	4
2.2	Weight Maps	4
3	Data Reduction	5
3.1	Calibration	5
3.2	Cross-Correlation Alignment	5
3.3	Cosmic Ray Rejection	6
3.4	Absolute Astrometry	7
3.5	Multidrizzle / Mosaic Combination	7
4	Data Quality	8
4.1	Exposure Depth and Sensitivity	8
5	Terms of Use	8
6	Contact Information	8

1 General Information

This document presents data from the Cosmic Assembly Near-infrared Deep Extragalactic Legacy Survey, CANDELS (Grogin et al. 2011; Koekemoer et al. 2011), PI: S. Faber, Co-PI: H. Ferguson, which is a 902-orbit Multi-Cycle Treasury program on the *Hubble Space Telescope*, designed to document the first third of galactic evolution from redshift $z \sim 1.5$ to 8 via deep imaging of more than 250,000 galaxies, as well as measuring Type Ia SNe beyond $z > 1.5$. The project targets five premier existing survey fields on the sky, each with extensive multi-wavelength data, and obtains new observations using the Wide Field Camera 3 (WFC3) IR and UVIS channels, together with the Advanced Camera for Surveys (ACS), in a variety of filters in the optical and near-infrared bands. Two different mini-surveys are included: the CANDELS/Deep survey covering ~ 130 square arcminutes within GOODS-N and GOODS-S and the CANDELS/Wide survey covering a total of ~ 720 square arcminutes in GOODS and in three additional fields (COSMOS, EGS, and UKIDSS/UDS).

Further details are in Koekemoer et al. (2011) for the data products, and Grogin et al. (2011) for the survey and observational design; in this README file we summarize the basic properties of the HST data products. Please also visit the CANDELS Survey Project website at:

<http://candels.ucolick.org>

where additional information can be found including field maps, observing products, and other ancillary products. The CANDELS data are all accessible from the Multimission Archive at STScI (MAST) High-Level Science Product pages, at:

<http://archive.stsci.edu/prepds/candels>

In Table 1 we provide a brief summary of the different components of the survey, including the wide and deep portions as well as the filter coverage on each of these.

Table 1: CANDELS Survey Fields: Observational Summary

Field	Prog. ID	Cycles	WFC3/IR Filters	WFC3/UVIS Filters	ACS/WFC Filters	#Exp.
UKIDSS/UDS (Wide)	12064	18	F125W,F160W	F350LP	F606W,F814W	752
GOODS-S Deep	12061,2	18,19	F125W,F160W	F350LP	F606W,F814W,F850LP	1740
GOODS-S Wide	12060	18,19	F125W,F160W,F105W	F350LP	F606W,F814W,F850LP	268
EGS (Wide)	12063	18,20	F125W,F160W	F350LP	F606W,F814W	756
GOODS-N Deep	12442,3	19,20	F125W,F160W,F105W	F350LP	F606W,F814W,F850LP	1800
GOODS-N Wide	12444,5	19,20	F125W,F160W,F105W	F350LP,F275W,F336W	F606W,F814W,F850LP	520
COSMOS (Wide)	12440	20	F125W,F160W	F350LP	F606W,F814W	752
Total						6588

The observations for each field are divided into several different epochs to facilitate searching for supernovae, with the deep fields (on GOODS-N and GOODS-S) containing a total of 10 separate epochs, while the wide fields (COSMOS, EGS and UDS and the GOODS wider areas) contain two epochs each. In addition, in some cases only a portion of the field could be observed in a given epoch, while another portion is observed separately (for example, EGS needed to be broken into two pieces in this way). Finally, the F105W imaging is obtained separately in some cases. A detailed summary of the observing schedule for all the fields, broken down by date, is available here:

http://candels.ucolick.org/survey/Obs_Schedule.html

2 Data Products

2.1 Science Images

The science images are all combined using `MultiDrizzle` (Koekemoer et al. 2002), which uses the `Drizzle` program (Fruchter & Hook 2002) to produce an image consisting of the weighted sum of all the input exposures. Further details about the combination algorithms are presented in these papers. The images are calibrated in units of electrons/s, with a pixel scale of 60 mas/pixel for the WFC3/IR data and 30 mas/pixel for the ACS/WFC and WFC3/UVIS data, and use a square kernel with a `pixfrac` setting of 0.8, well matched to the number of exposures and the dither strategy. The images have been astrometrically registered to an external catalog of the field, which in the case of the UDS consist of the WFCAM K-band catalog (Almaini et al. 2011, in prep.). Here we describe further details about the images, the processing, and additional data products. For additional details, see Koekemoer et al. (2011) for the HST data processing, and Grogin et al. (2011) for details about the observational planning and design.

2.2 Weight Maps

When combining images where each pixel value has noise associated with it, the optimal approach involves weighting each pixel according to the inverse square of its noise. In particular, it can be demonstrated that the appropriate quantity to use consists of the background noise only, and not including the noise associated with the emission from the object since that would lead to a biased estimation of the true flux in the pixel and result in photometric errors. Our CANDELS pipeline constructs inverse variance images for all the exposures obtained with WFC3/IR, WFC3/UVIS and ACS/WFC, using essentially the same formalism as described in Casertano et al. (2000) for the original WFPC2 Hubble Deep Field South observations, but modified to take into account the differences in the data formats for the ACS and WFC3 detectors. In particular, these detectors are calibrated to electrons per second and have been corrected for the detector gain, while the WFPC2 data are in DN and therefore still need to have the gain included in the calculation. Specifically, the WFC3/IR images are in electrons per second while WFC3/UVIS and ACS/WFC are in units of total electrons; however the input inverse variance images for all of them are calculated to be in units of counts per second, since the ACS images are scaled to be in electrons/second while drizzling and since that is also what the output units of the mosaics are in all cases. Given this, the formula that is used to calculate the inverse variance images for all three cameras, in units corresponding to inverse (electrons/second)² (including also a separate scaling by $scale^4$) is given by:

$$Inverse\ Var. = \frac{(ft)^2}{(D + fB) + \sigma_{read}^2} \quad (1)$$

where f is the inverse flatfield (as defined in the conventions for the flatfield reference files used in calibration), t is the exposure time (in seconds), D is the total accumulated dark current signal during the exposure, B is the total accumulated background level during the exposure, and σ_{read} is the read-out noise, with all three of the latter quantities being in units of electrons.

3 Data Reduction

3.1 Calibration

Each of the ACS/WFC and WFC3/UVIS exposures are initially calibrated with the Pyraf/STSDAS tasks `calacs` and `calwf3` respectively, for which the standard steps are quite similar, consisting of overscan / bias and dark correction, flatfielding, gain correction, and photometric calibration. In addition, after standard calibration, for the ACS/WFC images we applied two more corrections. The first one is the correction of bias striping that are the consequence of $1/f$ noise in the new detector electronics for ACS (Grogin et al. 2010). The second correction involves an algorithm to “de-trail” the deferred-charge trails that are related to charge transfer efficiency (CTE) degradation on the ACS detectors (Anderson & Bedin 2010). This results in a set of exposures for WFC3/UVIS and ACS/WFC that are significantly improved compared with the standard pipeline calibration.

The raw WFC3/IR raw exposures, still containing all the separate MULTIACCUM read samples, are calibrated using the Pyraf/STSDAS task `calwf3`. This task populates the bad pixel arrays using known bad pixel tables, followed by a bias subtraction for each read, based on the mean value of all the pixels in the 5-pixel wide reference pixel region. It then carries out a subtraction of the zeroth read in order to remove the bias structure across the detector, followed by a subtraction of the dark current reference files for either the SPARS50 or SPARS100 readout sequences, as applicable. This is followed by the non-linearity correction and photometric keyword calculation, using the current filter throughput tables and detector quantum efficiency curves. We also applied additional bad pixel masking, based on identification of hot pixels that were evident across multiple exposures.

Once all the separate MULTIACCUM reads are calibrated as described above, they then move through the next steps in `calwf3`, namely up-the-ramp slope fitting and cosmic ray rejection. For each pixel, this step performs a linear fit to the accumulating counts that are sampled during each MULTIACCUM read, while rejecting outliers from the fit as being due to cosmic rays. Currently a threshold of 4 sigma is used for this rejection, and flags are populated in the data quality arrays corresponding to the read during which the cosmic ray occurred. A final count-rate value is then computed for each pixel using only the unflagged reads, and is stored at the count-rate in the final calibrated exposure, while the uncertainty in the slope of counts versus time is stored in the error extension of the image. The final steps of `calwf3` includes multiplicative corrections for the detector gain and the flatfield structure across the detector, appropriate to each filter.

3.2 Cross-Correlation Alignment

To mitigate the uncertainties in the relative astrometric alignment between all the exposures in each visit, a cross-correlation technique is applied to all the exposures for a given camera in each visit. This consists of first passing all the exposures for each camera (including all filters) through a partial run of `Multidrizzle`, up to the point where a single-drizzled image is produced for each separate exposure. These images are then masked to eliminate cosmic rays and spurious sources, retaining non-zero pixel values only around objects that are large enough to provide sufficient signal for the cross-correlation. The area containing each object is tapered so as to avoid introducing artificial ringing in the cross-correlation step, and bright saturated sources, in particular stars with long charge-bleed columns are also eliminated since such columns in fact degrade the quality of the cross-correlation solutions.

The cross-correlation is then run separately for each of the different instruments in each orbit, but including all the filters. This is motivated by the fact that, despite some differences in the structure of galaxies between filters, these tend to randomize out and therefore have no net correlated signal for all the sources over an image, while a correlated signal such as a residual pixel shift remains much more dominant. After a set of first-pass cross-correlation shifts has been obtained, these shifts are then propagated back to the input exposures, in order to do a second-pass drizzle which includes a cosmic-ray rejection step with the improved shifts. The pixels flagged as cosmic rays in each exposure then have their flux replaced with pixels from the clean drizzled image, then subsequently re-drizzled to a new set of single-drizzled images for each separate exposure, which are again masked and used as input for the second-pass cross-correlation step. The cross-correlation shifts during the second-pass are typically much smaller than those measured during the first, indicating that the bulk of the shifts have already been accounted for, with the remaining shifts being primarily to correct small residuals that are introduced by cosmic rays in the galaxies used during the first pass. The second-pass shift corrections are also propagated back to the input exposures and are stored as the final set of relative coordinates for all the exposures within a given orbit.

3.3 Cosmic Ray Rejection

With the relative shifts within each orbit having been corrected, the next step consists of creating a cosmic ray mask for all the exposures of a given filter, for each camera, during a given orbit, by carrying out another run of `MultiDrizzle`, this time with the improved relative shifts. The cosmic rays are identified in the `driz_cr` step of `MultiDrizzle` using a process that first creates a series of separately drizzled images, one for each input exposure. These are subsequently used to create a median image using the “minmed” algorithm in `MultiDrizzle`, which enables the minimum to be used instead of the median in cases where valid pixels from only 2 or 3 exposures are present, if one of them exceeds the others by $> 4\sigma$. The clean median image is then transformed back to the distorted detector frame of each input exposure to carry out cosmic ray rejection using the following approach. The input counts in a given pixel in the original exposure, I_{exp} are compared with the counts from the median image I_{med} for the same pixel, together with the derivative of the median image Δ_{med} , defined as the steepest gradient from that pixel to its surrounding pixels (with all these quantities being in units of electrons). A pixel is flagged as a cosmic ray if it exceeds a threshold: $|I_{\text{exp}} - I_{\text{med}}| > S\Delta_{\text{med}} + SNR\sqrt{\sigma_{\text{read}}^2 + |I_{\text{med}} + B|}$ where S and SNR are adjustable scaling factors, and B is the background sky value that was measured for the exposure. The inclusion of the gradient term Δ_{med} effectively “softens” the cosmic ray rejection in regions of steep gradients such as bright cores of objects, where pixel-to-pixel variations can exceed simple Poisson statistics.

For the WFC3/IR images, most of the cosmic rays are already rejected during the up-the-ramp sampling, and are excluded from the process of fitting a slope to the measured counts when determining the count-rate in each pixel. However, there are occasional cosmic rays that are not fully removed, or warm pixels that are not accounted for in the dark file correction, so the WFC3 exposures for each filter are also passed through this step. For the ACS/WFC data, there are typically between 2 and 4 exposures per filter in a given orbit. Given the typical cosmic ray rate of $\sim 1 - 2\%$ during our exposure times, this means that for a 4-exposure depth, $\lesssim 1 - 2$ pixels can be expected to be hit by cosmic rays during all 4 exposure, while for 2 exposures this number increases to $\sim 2000 - 6000$ pixels would be affected by cosmic rays during both exposures. However, this is still only $\sim 0.01 - 0.04\%$, meaning that ~ 1 out of every 100 small galaxies (~ 100 pixels in area) would be affected, losing $\sim 1 - 4\%$ of its pixels on average.

3.4 Absolute Astrometry

For all the exposures of a given instrument in each visit, an “astrometric detection image” is then produced by doing another run of `MultiDrizzle`, to produce a single combined image containing all the exposures of all the filters for that instrument, applying the cosmic ray masks that have been produced when combining the filters separately. This is motivated by the fact that the relative shifts for all the exposures, for all filters in a given orbit, have now already been determined to an accuracy of $\sim 0.5 - 1$ mas, therefore nothing further would be gained by now attempting to solve for shifts separately for different filters in a given orbit. Therefore, the exposures for all the filters are combined into a single image (separately for each different instrument), yielding one image for WFC3/IR, another for WFC3/UVIS, and another for ACS/WFC, for each visit. A catalog is then produced from this single multi-filter image, which also has the advantage of providing increased depth and reducing the impact from cosmic rays compared with the individual filter images.

All the sources in the multi-filter catalog for each visit are then matched to the sources in the relevant portion of the external catalog, using a number of iterative steps. The first iteration uses a relatively large tolerance (up to a few arcseconds) and only the brightest $\sim 20 - 30$ sources in each image, in order to determine the dominant terms in the shifts for right ascension and declination. Once these have been accounted for, several additional iterations are carried out using the full catalog of sources in each image, using progressively tighter matching tolerances down to $0''.1$, and solving for the residual remaining shifts as well as the rotation errors due to the uncertainties in guidestar position. For all visits, $\sim 300 - 400$ sources are typically then matched at the faintest levels and tightest tolerances between the HST `Multidrizzle`-combined image and the reference catalog.

3.5 Multidrizzle / Mosaic Combination

In preparation for the final step of combining all the exposures for each filter into a single mosaic using `MultiDrizzle`, our pipelines first create for each exposure a corresponding inverse variance image, which contains all the “intrinsic” error terms associated with each pixel (including noise from accumulated dark current, detector read-out, and photon noise from the background as modulated multiplicatively by the flatfield and the detector gain), but not the additional “extrinsic” Poisson noise from the astronomical sources in the image. This approach was described by Casertano et al. (2000) for the Hubble Deep Field South, and has also been used in GOODS (Giavalisco et al. 2004), COSMOS (Koekemoer et al. 2007; Scoville et al. 2007), AEGIS (Davis et al. 2007), UDF (Beckwith et al. 2006) and other projects, and is implemented as a routine option in `Multidrizzle`.

The final pass of `MultiDrizzle` is then run using the above output pixel scale and `pixfrac` settings, and applying the inverse variance weight image associated with each exposure which contains the full set of masks from cosmic rays, bad pixels, satellite trails and other blemishes in the detector. In each case the images for all filters are drizzled onto a common tangent plane projection on the sky, which is defined to match the existing ones where known, to facilitate a direct comparison with pre-existing data on these fields. Four of the fields (GOODS-North, GOODS-South, COSMOS and EGS) all have an existing tangent plane point already defined which we adopt for the CANDELS mosaics as well. For the UDS field, we adopted a common tangent plane point designed to satisfy the current surveys on that field.

4 Data Quality

4.1 Exposure Depth and Sensitivity

The prime WFC3 exposure time consisted of F350LP, F125W and F160W imaging, with exposure times as shown below. In parallel, ACS/WFC exposures were obtained in the F606W and F814W filters. Most pointings were covered by ACS parallels which provided the optical baseline for the supernova search, except for the easternmost WFC3/IR exposures, which were instead covered by short WFC3/UVIS exposures. Since the WFC3 exposures are contiguous, this means that the ACS exposures overlap each other (due to the larger size of ACS/WFC), thus effectively doubling the exposure depth. The resulting sensitivities corresponding to these exposure times are shown below, quoting 5σ point-source sensitivities in AB magnitudes.

Table 2: Data Quality Summary

Instrument	Filter	Eff. Exptime	Sensitivity	PSF FWHM
ACS/WFC	F606W	3300s	28.5	0''08
ACS/WFC	F814W	6900s	28.4	0''09
WFC3/UVIS	F350LP	434s	26.4	0''08
WFC3/IR	F125W	1900s	27.0	0''12
WFC3/IR	F160W	3200s	26.9	0''18

5 Terms of Use

When referencing the CANDELS project and data products, please cite both these papers together:

- [Grogin, N. A. et al., 2011 ApJS, 197, 35](#) (“CANDELS Survey and Observational Design”)
- [Koekemoer, A. M. et al., 2011, ApJS, 197, 36](#) (“CANDELS HST Imaging Data Products”)

Also, when using data from the CANDELS project, please include the following acknowledgement:

“This work is based on observations taken by the CANDELS Multi-Cycle Treasury Program with the NASA/ESA HST, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS5-26555.”

6 Contact Information

For any questions about this dataset please contact:

Anton M. Koekemoer, STScI (<http://www.stsci.edu/~koekemoer>)

Further details are also available at the CANDELS Survey Project website:

<http://candels.ucolick.org>

and the Multimission Archive at STScI (MAST) High-Level Science Product pages:

<http://archive.stsci.edu/prepds/candels>

References

- Anderson, J. & Bedin, L. 2010, *PASP*, 122, 1035
Beckwith, S. V. W. et al. 2006, *AJ*, 132, 1729
Casertano, S. et al. 2000, *AJ*, 120, 2747
Davis, M. et al. 2007, *ApJL*, 660, 1
Fruchter, A. S. & Hook, R. N. 2002, *PASP*, 114, 144
Giavalisco, M. et al. 2004, *ApJL*, 600, L93
Grogin, N. A. et al. 2010, *HST Calibration Workshop*
Grogin, N. A. et al. . 2011, *ApJS*, 197, 35
Koekemoer, A. M. et al. 2011, *ApJS*, 197, 36
Koekemoer, A. M. et al. 2007, *ApJS*, 172, 196
Koekemoer, A. M. et al. 2002, *HST Calibration Workshop*, 337
Scoville, N. Z. et al. 2007, *ApJS*, 172, 1