Instrument Science Report ACS 2022-03
Instrument Science Report WFC3 2022-06

Improved Absolute Astrometry for ACS and WFC3 Data Products

J. Mack, W. Hack, M. Burger, R. L. White, V. Bajaj,
R. J. Avila, G. S. Anand, M. de la Pena

August 23, 2022

ABSTRACT

As of late-2019, MAST data products for ACS and WFC3 include improved absolute astrometry in the image header World Coordinate System (WCS). The updated WCS solutions are computed during pipeline processing by aligning sources in the HST images to a select set of reference catalogs (e.g. Gaia eDR3). We compute statistics on the alignment fraction for each detector and estimate the uncertainties in the WCS solutions when aligning to different reference catalogs. We describe two new types of Hubble Advanced Products (HAP), referred to as Single Visit Mosaics (SVMs) and Multi Visit Mosaics (MVM), which began production in MAST in late-2020 and mid-2022, respectively. The SVM products include an additional relative alignment across filters in a visit, and the drizzled images are used to generate point source and segment catalogs during pipeline processing. These catalogs supersede those produced by the Hubble Legacy Archive and will be the basis of the next version of the Hubble Source Catalog. The MVM data products combine all ACS/WFC, WFC3/UVIS, or WFC3/IR images falling within a pre-defined 0.2° x 0.2° 'sky cell' for each detector+filter, which are drizzled to a common all-sky pixel grid. When combining observations over a large date range, MVMs may have photometric errors of several percent or systematic alignment errors when combining visits with different catalog solutions. We therefore recommend these to be used as 'discovery images' for comparing observations in different detectors and passbands and not for precise photometry.

1. Introduction

The header of each calibrated ACS and WFC3 two-dimensional FITS image contains information to translate the pixel positions of celestial objects into sky coordinates in RA and Dec. This information relies on an accurate knowledge of the optical distortions in each HST detector to rectify the images prior to alignment (see Kozhurina-Platais et al. 2015 for ACS/WFC, for example). The rectified pixel positions may then be translated into sky positions using the image header 'World Coordinate System' (WCS), based on coordinates in the Guide Star Catalog (GSC) used at the time of the observation. For a detailed description of astrometric information carried in the FITS image headers, see Chapter 4 of the DrizzlePac Handbook (Hoffmann et al. 2021).

The GSC was constructed to support the pointing and target acquisition of HST and was created by digitizing photographic plates taken from ground-based telescopes (Lasker et al. 1990). The first version used by HST operations (GSC v1.1) had nominal RMS errors of 500 mas per coordinate, with errors as large as ~1-3 arcsec near the plate edges. As the accuracy of the GSC has been refined over time, the pointing accuracy of HST has improved as a result. With the introduction of GSC v2.3 in October 2005 (during Cycle 15), the RMS error per coordinate was reduced to ~300 mas over the whole sky. The latest GSC v2.4.0 was released in October 2017, improving the celestial coordinates with positions from Gaia DR1 and reducing errors to < 30 mas over the entire sky. Including small ~100 mas errors in the alignment of the science instruments to the HST focal plane, the total error in the absolute astrometry is ~200 mas when using GSC v2.4.0. For a summary of 'Key Guide Star Catalog releases and associated errors', see Section 4.5.1 of the Drizzle Handbook (Hoffmann et al. 2021).

The primary goal of the 'HST Astrometry Project' at STScI was to correct any pointing errors in archival ACS and WFC3 data products, when possible, by aligning the distortion-corrected frames to one of three reference catalogs (Gaia eDR3, GSC v2.4.2, or 2MASS). A secondary goal was for the MAST archive to automatically generate 'Hubble Advanced Products' (HAP), including both mosaicked images and source lists which can be accessed from the standard MAST interfaces. These new products are intended to replicate the advanced products from the Hubble Legacy Archive (HLA), including those used to construct the Hubble Source Catalog (Whitmore et al. 2016). They further extend the mosaicking process by combining all HST observations in a given region of the sky for each combination of detector+filter, referred to as 'layers'. These are all drizzled to a pre-defined, all-sky pixel grid, allowing for convenient way of comparing observations over large range of wavelengths and passbands.

In this report, we describe the set of reprocessed HST observations in Section 2 and the methodology used for alignment in Section 3. We describe the software for generating 'advanced' products in Section 4, and we provide recommended use cases for the three 'levels' of MAST products. In Section 5, we highlight the alignment success fraction for three reference catalogs, including the estimated errors for each detector as a function of date and the number of catalog matches. In Section 6, we describe file naming conventions for the three levels of MAST products and helpful FITS keywords which may be used to assess the quality of the alignment. We describe new source catalogs which are generated from SVM products in Section 7, and we highlight software tools for working with the updated ACS and WFC3 data products in Section 8. In Appendix A, we choose a sample HST sky cell and list all files associated with the MVM, SVM, and standard data products. Finally, we showcase several examples of the MVM color mosaics in Appendix B, which were created from calibrated images from different detectors and filters and drizzled to a common all-sky grid.

2. Observations

All ACS and WFC3 external observations were reprocessed in MAST in order to update the image header WCS and to produce high-level 'advanced' data products. The Advanced Camera for Surveys was installed in March 2002 in Servicing Mission 3B, while the Wide Field Camera 3 was installed in May 2009 in Servicing Mission 4. The total number of external exposures available for download from MAST as of April 2022 are reported in Table 1. Only a fraction of these can be aligned to catalogs, and statistics on the alignment results are reported for each detector. These statistics do not include subarray observations, which typically have a lower alignment fraction than full-frame images due to their smaller field of view. (A large number of these are calibration images of a single star and cannot be directly aligned.)

Additionally, some observation modes are *excluded* from catalog alignment due to the complexity of the data and/or lack of point sources for matching. These modes include:

Prisms: ACS/(HRC, SBC), **Grisms**: ACS/(WFC, HRC), WFC3/(UVIS, IR),
Ramp filters: ACS/WFC, **Quad filters**: WFC3/UVIS,
Polarizers: ACS/(WFC, HRC), **Coronagraph:** ACS/HRC,
Earth flats: ACS/(WFC, HRC), WFC3/(UVIS, IR),
Moving targets: all detectors, **Drift-And-SHift programs** (WFC3/IR)

While the alignment may be unsuccessful, either due to a guide star failure or a lack of sources in the HST image or the reference catalog, we find that approximately **90% of ACS/WFC, 70% of WFC3/IR and 60% of WFC3/UVIS** *alignable* datasets have been successfully aligned to any catalog (see Table 1, column 7). On the other hand, only **20% of ACS/HRC** images and ~**10% of ACS/SBC** images are successfully aligned. For further discussion, see Section 5.

Table 1: Alignment statistics for WFC3 and ACS images through April 2022. Columns 2-5 give the range of observation dates for each detector, the number of reprocessed images (to the nearest hundred), and the number and fraction of *alignable* images. Columns 6-8 list the number (and corresponding fraction) of successfully aligned images and the fraction of *alignable* images that were fit to Gaia eDR3.

Detector	Date Range	Number Datasets	Number Alignable	Fraction Alignable	Number Aligned	Fraction Aligned	Fraction eDR3
WFC3/UVIS	2009 - 2022	107,000	51,500	0.48	31,000	**0.60**	0.53
WFC3/IR	2009 - 2022	141,800	61,600	0.43	42,400	**0.69**	0.58
ACS/WFC	2002 - 2022	106,000	93,200	0.88	84,200	**0.90**	0.79
ACS/HRC	2002 - 2007	40,000	25,400	0.63	5,000	**0.20**	0.12
ACS/SBC	2002 - 2022	16,200	12,800	0.79	1,400	**0.11**	0.09

3. Alignment

3.1 Methodology

On December 3, 2019, MAST released the first set of ACS and WFC3 observations with improved absolute astrometry. As observations are processed and/or reprocessed in the pipeline, their WCS is updated to use the most accurate solution available, with two levels of correction possible:

- ***a priori*** *:* corrects the coordinates of the guide stars in use at the time of observation to their coordinates as determined by Gaia DR1, applying a global offset to the WCS
- ***a posteriori*** : cross-matches sources in the HST image with positions from an external reference catalog to improve the WCS by fitting (x/y) to (RA/Dec).

By combining the updated guide star coordinates with knowledge of the instrument distortions, the *a priori* correction reduces the pointing uncertainties to ~200 mas over the entire sky. (Only observations executed prior to October 2017 will require an *a priori* correction, e.g. when the HST pointing was based on catalogs prior to GSC v2.4.0.)

When possible, an additional *a posteriori* correction is computed by aligning sources in the HST images to one of the three reference catalogs. These corrections are *only* computed for images with an adequate number sources matched (see Table 2). They remove uncertainties in the focal plane alignment and are expected to have the smallest absolute astrometric errors.

The software used to align and drizzle MAST data products is described on the `DrizzlePac` 'Pipeline Astrometric Calibration' webpage[1]. The alignment portion of the code will loop over catalogs, fit 'methods', and fit 'geometries', using the first successful fit it computes, with:

- catalogs: **GAIAeDR3**, **GSC242**, or **2MASS** (in this order)
- methods: **image-by-image, relative,** or **single visit**
- geometries: **rscale**, **rshift**, or **shift** (depending on the number of cross-matches)

The fitting between the matched source lists (HST to catalog) is performed by the `TWEAKWCS` module `linearfit`[2], where the parameter '`fitgeom`' describes the geometry used for computing offsets, rotations and/or scale changes. Table 2 provides a summary of the fitting geometry used by the software as a function of the number of successfully matched sources. Note that in order to preserve the integrity of the distortion calibration for each detector, these transformations do not allow for higher order solutions, such as the 'general' fit which allows for skew terms (e.g. an independent scale and rotation for each axis).

[1] https://drizzlepac.readthedocs.io/en/latest/runastrodriz.html
[2] https://tweakwcs.readthedocs.io/en/latest/source/linearfit.html

Table 2: Fitting geometry used for different numbers of cross-matches between the HST source list and the reference catalog.

Fit Geometry	Nmatches	Description
'rscale'	$N \geq 6$	X and Y shift, global rotation, and scale
'rshift'	N = 4, 5	X and Y shift, global rotation
'shift'	N = 3	X and Y shift only
---	$N \leq 2$	No fit

When Gaia eDR3 was released in December 2020 (Gaia Collaboration, Brown, A.G.A., et al., 2021), all ACS and WFC3 datasets were reprocessed in MAST, and nearly every prior Gaia DR2 or DR1 WCS solution was replaced with Gaia eDR3. The updated Gaia catalog is based on 34 months of observations, compared to 22 months for DR2, and has more accurate proper motion information. While the latest Gaia DR3 catalog (released in June 2022) will be used for alignment at a future date, the astrometry is identical to eDR3 so the HST alignment results are not expected to change.

When no successful Gaia solution is found, the software then attempts to align to GSC v2.4.2, which is based on the 'Guide Star Catalog II plus PanSTARRS, SMSS, WISE, GALEX, and ultraVISTA, with astrometric data updated to match Gaia DR2' (Lasker et al. 2021). The third catalog used for alignment is 2MASS (Skrutskie et al. 2006), which has astrometric accuracy of order 100 mas and is based on observations from two 1.3-meter diameter telescopes at Mount Hopkins, Arizona, and Cerro Tololo, Chile.

3.2 WCS Solutions

The active WCS solution may be identified via the science extension header keyword `WCSNAME` which follows the convention: **WCSNAME = 'Original Solution–Correction Type'**,

where **Original Solution** may be either:

- OPUS : ground system WCS, no distortion correction
- IDC_* : distortion-corrected WCS, where * is the name of the IDCTAB reference file

and **Correction Type** may be either an **'*a priori*' WCS** solution:

- GSC240 : GS coordinates are corrected from the original frame (GSC v1.1, GSC v2.3) to the Guide Star Catalog v2.4.0 frame, which is based on Gaia DR1
- HSC30 : GS coordinates are corrected to the Hubble Source Catalog v3.0 frame, which is based on Gaia DR1

or an **'*a posteriori*' WCS** solution:

- FIT-IMG-cat : each image is separately aligned to the reference catalog
- FIT-REL-cat : all images in the same filter are aligned per visit, before a global catalog fit
- FIT-SVM-cat : all images in every filter are aligned per visit, before a global catalog fit

Successfully aligning to a reference catalog will result in an *a posteriori* solution designated with `WCSNAME` = '`IDC_*-FIT-*-cat`', where the '`*`' represents the fit method listed in Section 3.1. A new 'headerlet' extension will also be added to the FITS file which includes specific WCS keywords defining the transformation from HST pixels to sky coordinates (RA, Dec). As an HST dataset is reprocessed in MAST, additional headerlets may be added to the FITS image, with the *best* solution applied to the primary (active) WCS. For more detail on the use of headerlets, see the `DrizzlePac` 'Astrometry in Drizzled Products' webpage[3].

Other diagnostic keywords are populated to the image headers, and these allow the user to inspect the quality of the fit for each individual dataset. These keywords include the number of catalog cross-matches (`NMATCHES`), the fitting geometry used for the transformation (`FITGEOM`), and the RMS of the fit residuals in milli-arcseconds for both RA and Dec coordinates (`RMS_RA`,`RMS_Dec`). These keywords are recorded in the header of the calibrated and drizzled data products so that analysis may be performed on either, depending on the science goals. These plus additional keywords specific to HAP drizzled data products are described in Section 6.

4. Advanced Products

4.1 Single Visit Mosaics

On December 17, 2020, MAST began production of advanced ACS and WFC3 data products. These Hubble Legacy Archive (HLA)-style mosaics are collectively aligned to a common astrometric reference frame and are referred to as Single Visit Mosaics (SVMs), since they comprise all of the observations from a single HST visit. (Note that a visit is defined as a series of exposures of a given target that execute in one or more consecutive orbits using the same set of guide stars.) SVMs have an additional 'relative' alignment across all filters in a visit by cross-matching sources in the HST images themselves. Once the relative alignment has been improved, the entire visit is globally realigned to the reference catalog. For example, SVM products with both relative alignment (by filter) and absolute alignment to Gaia eDR3 will contain the string `'FIT_SVM_GAIAeDR3'` in the `WCSNAME` keyword. For details on SVM processing, see the `DrizzlePac` 'Single Visit' webpage[4].

In contrast with standard HST products, which are drizzled using the native detector orientation at the time of observation, SVMs are drizzled with North oriented up. In some cases, this may introduce additional *correlated noi*se as a result of resampling when drizzle divides the power from a given input pixel between several output pixels. In general, the correlation between pixels varies across the detector field of view and depends on the choice of drizzle parameters, as well as the geometry and orientation of the dither pattern. For a detailed discussion of correlated noise, see Section 3.3.2 of the DrizzlePac Handbook (Hoffmann et al. 2021).

SVM products include the exposure-level (single) drizzled images, the combined images for each filter, and a total image which combines all exposures in the visit (see Section 5, Table 3). The

[3] : https://drizzlepac.readthedocs.io/en/latest/astrometry.html
[4] : https://drizzlepac.readthedocs.io/en/latest/singlevisit.html

total 'white light' image is used as a detection image for generating the point and segment catalogs each filter. MAST began providing source catalogs as part of the SVM data products on November 23, 2021. For more details, see Section 7.

The SVM drizzled images combine all images in a given filter, regardless of exposure time or association membership. As for standard HST products, the input images are weighted by the total exposure time of each pixel (`final_wht_type = 'EXP'`) when combining observations, and the weight extension of the drizzled image can be considered an effective exposure time map. Exposure weighting does not work well for saturated stars when the input frames have widely different exposure times. The problem occurs at the boundary ring of the two PSFs where `AstroDrizzle` uses the DQ flags to fill in saturated pixels in the long exposure with unsaturated pixels from the short exposure. `AstroDrizzle` will underweight the flux from the short exposure compared to the long one, causing a discontinuity in the PSF. This effect can make it difficult to model the PSF since the radial profile is compromised, and it will also make the flux too low in those boundary pixels. Thus, we recommend using the shorter exposure-level drizzled images for accurate photometry of very bright stars.

4.2 Multi-Visit Mosaics

On May 26, 2022, MAST began producing a second type of advanced products referred to as Multi-Visit Mosaics (MVMs). These combine all available public observations of ACS and WFC3 into separate layers by detector and filter, which are then drizzled to a common, pre-defined pixel grid over the whole sky. The MVMs adopt the WCS solutions from the SVM products as input (with no additional relative alignment). For ACS/WFC and WFC3/UVIS, only the CTE-corrected calibrated images are used to generate the drizzled mosaics, as done for SVMs.

MVM products rely on tessellation of the entire sky and use the same tiling system defined by the PanSTARRS project, as described on the 'PS1 Sky tessellation patterns' webpage[5]. Projection cell centers are located on lines of constant declination and spaced 4° apart. MVM processing extends each projection cell to 4.2° x 4.2° in size in order to provide sufficient overlap to allow data from one projection cell to be merged with data from a neighboring cell, if needed. These cells are indexed starting at 0 at $\delta = -90°$ (South Pole) and ending at 2643 at $\delta = +90°$ (North Pole).

Each projection cell is split into a grid of 20 x 20 'sky cells' which are spaced 0.2° apart. The sky cells are oversized at 0.244° x 0.244° (21953 x 21953 pixels) at a scale of 0.04″/pixel to allow for overlap with adjacent cells. Each sky cell is identified by its position within the projection cell with (x, y) coordinates ranging from 1-20, starting in the lower left corner. Because the WCS for each sky cell is defined as a subarray of the projection cell, observations spanning multiple sky cells share the same tangent plane projection, where the tangent point in RA and Dec is given by the image header keywords `CRVAL1` and `CRVAL2`. This means that targets spanning multiple sky cells may simply be stitched together without having to re-drizzle the observations. MAST provides an API for requesting custom cutouts of MVMs, including support for objects spanning multiple sky cells. For details, see Section 8.

[5] https://outerspace.stsci.edu/display/PANSTARRS/PS1+Sky+tessellation+patterns

Currently, all public observations within a sky cell for a given detector+filter are combined to form a layer, and MAST will regenerate the MVM layers as new observations or improved reference files become available or as proprietary data becomes public. The full sky cell is rarely saved; only the smallest rectangular section containing all pixels used by all layers in the mosaic is archived. For details on MVM processing, see the `DrizzlePac` 'Multi Visit Mosaic' webpage[6].

Initial testing shows that some visits in the same sky cell have global offsets of several pixels, and these will result in slightly smeared drizzled data products when combining images with different types of WCS solutions (e.g. *a priori* with *a posteriori*). For this reason, `AstroDrizzle` processing for MVMs does not perform any additional CR-flagging but instead uses the DQ flags (bit value = 4096) populated in the SVM calibrated data. Future enhancements to the MVM products could include enforcing all input datasets to share the same `WCSNAME` (e.g. '`FIT-SVM_GAIAeDR3`') across the entire projection cell. This change would most likely improve the quality of the alignment for a given layer, at the cost of reducing its total depth. Alternately, a separate layer including only datasets with Gaia eDR3 solutions could be created.

4.3 Use Cases for Different Products: HST, SVM, MVM

In Table 3, we compare the three levels of MAST data products and provide recommended use cases for each. The first column describes the level of product, where the MAST search parameter `PROJECT` is set to `'HST'` for standard pipeline products and `'HAP'` for advanced products. Recall that `HAP` includes both the SVM and MVM data products which are distinguishable by their different file names: `hst_prop_visit_*drz.fits` for SVMs and `hst_skycell-*_drz.fits` for MVMs (see Section 6). Column 2 lists the method used to align and combine datasets, while columns 3 and 4 provide considerations and recommended useage for each level.

5. Alignment Results

The `tweakwcs.linearfit` module performs a 2-D linear fit that transforms source positions in the reference catalog (in RA & Dec) to pixel positions on the detector (in x and y). The 'best' transform minimizes the fit residuals, after iteratively sigma-clipping outliers. The RMS of those residuals, hereafter referred to as the 'fit RMS', are recorded in the image header keywords `RMS_RA` and `RMS_Dec` in units of mas. Table 4 provides statistics on the alignment fraction and the fit RMS for standard '`HST`' products and for advanced '`HAP`' products. (MVMs adopt the SVM WCS solutions so are not separately listed). For images with no successful catalog alignment, or for modes which skip the alignment step, the WCS retains the best *'a priori'* solution which updates guide star coordinates to match either GSC v2.40 or HSC v3.0 (see Section 3.2).

[6] https://drizzlepac.readthedocs.io/en/latest/multivisit.html.

Table 3: Usage recommendations for standard MAST data products (Project= 'HST') and advanced data products (Project= 'HAP'), which includes both SVMs and MVMs. Column 2 describes the method for aligning and combining each set, column 3 highlights considerations for different observation modes, and column 4 provides some general usage guidelines.

DATA LEVEL	METHOD	CONSIDERATIONS	RECOMMENDED USE
HST	Aligns 'associated' and single exposures separately for each filter Combines only 'associated' datasets	WCS '*a posteriori*' solutions for multiple filters in a visit may have small systematic offsets Grism/direct images, blue/red filters may have different WCS solutions; If so, use '`updatewcs`' to reset the WCS to the *a priori* solution	Time-dependent photometry with < 1% errors Use short exposure drizzled products for photometry of saturated sources
HAP SVM	Improves the relative alignment of exposures in a given filter and of all filters in a visit Performs a global re-alignment of the visit Combines all datasets per filter per visit, regardless of association rules (e.g. custom dithers, short/long exposures, different IR sample-sequences)	Photometry of saturated sources may not be optimal in combined exposures (short and long) Drizzled products may not align with other visits, especially for large dithers which will have a different set of matched sources Drizzled with North oriented up so the files are up to √2 larger. Rotating may add correlated noise compared to standard HST products	Time-dependent photometry with < 1% errors Improved relative astrometry for multi-filter visits Grism/direct image pairs use the same *a priori* WCS Multi-filter aperture photometry provided via source catalogs
HAP MVM	Uses the SVM WCS for each visit, with no additional alignment Combines all exposures in a given sky cell per detector and filter, regardless of date or proposal ID Drizzles all detectors onto a common pixel grid at a scale of 0.04″/pixel for easy comparison	Drizzled files may be *very large* The combined layer may be degraded if visits with both *a priori* and *a posteriori* WCS solutions are combined. Proper motions can lead to poor quality images for some targets. Visits spanning a large range of dates do not account for time-dependent sensitivity, seen for some detectors Products do not account for differences in PSFs across detectors	'Discovery' images Photometry error <5% Ideal for comparing multiple detectors and filters across a large wavelength range Not produced for ACS HRC, ACS/SBC WFC3/IR images at 0.12″/pixel scale are much less noisy than the super-sampled data

Table 4: Statistics for *alignable* WFC3 and ACS observing modes, including the alignment fraction and the median fit RMS of the linear transformation from HST pixels to sky coordinates (see text for details).

Catalog	Detector	HST Fraction	HAP Fraction	HST RMS (mas)	HAP RMS (mas)	HST RMS (pix)	HAP RMS (pix)
eDR3	WFC3/UVIS	0.53	0.78	**10 ± 8**	**13 ± 10**	0.3	0.3
	WFC3/IR	0.58	0.53	**24 ± 15**	**32 ± 17**	0.2	0.3
	ACS/WFC	0.79	0.92	**16 ± 11**	**17 ± 13**	0.3	0.4
	ACS/HRC	0.12	0.24	**17 ± 8**	**21 ± 8**	0.6	0.8
	ACS/SBC	0.09	0.12	**30 ± 7**	**35 ± 7**	1.2	1.4
GSC242	WFC3/UVIS	0.05	0.08	**22 ± 13**	**41 ± 17**	0.6	1.0
	WFC3/IR	0.10	0.28	**111 ± 36**	**119 ± 25**	0.9	0.9
	ACS/WFC	0.07	0.05	**38 ± 39**	**56 ± 55**	0.8	1.1
	ACS/HRC	0.07	0.02	**20 ± 6**	**22 ± 8**	0.8	0.8
	ACS/SBC	0.01	0.02	**18 ± 6**	**37 ± 12**	0.7	1.5
2MASS	WFC3/UVIS	0.03	-	**51 ± 9**	-	1.3	-
	WFC3/IR	0.01	-	**92 ± 44**	-	0.7	-
	ACS/WFC	0.05	-	**57 ± 21**	-	1.1	-
	ACS/HRC	-	-	-	-	-	-
	ACS/SBC	-	-	-	-	-	-
'a priori' (No Fit)	WFC3/UVIS	0.40	0.15	-	-	-	-
	WFC3/IR	0.31	0.20	-	-	-	-
	ACS/WFC	0.10	0.03	-	-	-	-
	ACS/HRC	0.80	0.74	-	-	-	-
	ACS/SBC	0.89	0.86	-	-	-	-

In Table 4, we list the '`HST`' and '`HAP`' alignment fraction in columns 3 and 4 for four possible WCS solutions, where the four rows sum to 1.0 for each detector. The ACS/WFC detector has the highest success rate, with only 10% of HST and 3% of HAP images *not* receiving an *a posteriori* solution. Considering the ~20 year ACS observing timeline, the high success rate for images aligned to Gaia eDR3 (80% for HST and 92% for HAP data) is a result of including proper motion information and is a significant improvement on the results obtained with Gaia DR1. This earlier version of the catalog had both a lower alignment fraction and larger fit RMS values, due to the lack of proper motion data.

In columns 5 and 6, we provide an estimate of the alignment error for each detector, computed by adding the `RMS_RA` and `RMS_DEC` keyword values in quadrature and then taking the median over all datasets. The error bars reflect the spread in the fit RMS values each detector, computed as the standard deviation of the mean over all datasets, after clipping unusually large values. Columns 7 and 8 give the same median RMS value in pixels, computed using the native scale of the drizzled images for each detector (see Table 9, column 2). For ACS/WFC, WFC3/UVIS, and WFC3/IR, the median fit RMS is ~0.3 detector pixels for Gaia eDR3 WCS solutions, corresponding to ~0.2 pixels in each dimension (RA & Dec). Observations aligned to GSC242 and 2MASS have larger median fit RMS values of up to ~1 pixel, corresponding to ~0.7 pixels in each dimension.

The fit RMS for each image is plotted as a function of the number of catalog matches in the top panels of Figures 1 – 5 for each detector. The fit RMS is plotted as a function of observation date in the center and bottom panels of each figure for '`HST`' and '`HAP`' products, respectively, where the range of the y-axis range is reduced (compared to the top panel) to show RMS values more typical of a successful catalog alignment. The dashed horizontal line highlights a subset of images with very large RMS values, e.g. > 70 mas for WFC3/UVIS. A spot check of the drizzled products show that these typically indicate poor alignment, and we therefore exclude them when computing the RMS statistics in Table 4. WCS solutions with very large RMS values will be investigated for a future data release, and the active WCS may ultimately be reset to the *a priori* solution.

For WCS solutions based on fewer than 10 matched sources, the scatter in the fit RMS increases rapidly, with an envelope of both low and high values. These are especially pronounced for ACS/WFC and WFC3/UVIS images aligned to 2MASS and GSC242, which have larger positional uncertainties. These WCS solutions may contain small, systematic errors in the alignment due to the number of free parameters used for the fit geometry and should be treated with caution. Currently at least six matches (see Table 2) are required for the '`rscale`' geometry, which applies a shift, rotation, and scale, but this assumes a uniform (random) distribution of sources across the array.

In some cases, the fit RMS of a given image may be underestimated if the sources do not randomly sample the field of view, *e.g.* if they all fall within one corner of the detector or along a diagonal. In the case of very few catalog matches, an improved alignment across visits may be possible by modifying the fit requirements listed in Table 2 or by having a different set of requirements depending on the accuracy of the reference catalog positions.

5.1 Gaia eDR3

For each of the five detectors, the **most common catalog solution is Gaia eDR3**. As shown in Table 4, the success fraction for *alignable* datasets is 0.5 for WFC/UVIS, 0.6 for WFC3/IR, and 0.8 for ACS/WFC standard 'HST' data products. ACS/HRC and ACS/SBC images rarely have an *a posteriori* solution, and the success fraction for eDR3 alignment is only ~0.1 for each. These detectors have a much smaller field of view (and thus fewer available sources for alignment), and they suffer from a lack of UV components matching with those in the visible wavelength catalogs.

The **alignment fraction is typically higher for HAP versus HST** products, with the exception of WFC3/IR and ACS/SBC where it is roughly the same. The success fraction for HAP increases from 0.5 to 0.8 for WFC3/UVIS, 0.8 to 0.9 for ACS/WFC, and 0.1 to 0.2 for ACS/HRC. This is likely due to the improved relative alignment across filters in a visit for HAP data, before a global realignment. For example, the 42 UVIS filters contain a large number of UV and narrowband filters which are more difficult to separately align to Gaia, but more likely to align when treated as a set, using a combined source list for all filters in the visit. For ACS/HRC, the alignment success fraction doubles for HAP products, with similar RMS values as for HST products. For ACS/SBC, the success fraction is ~0.1 for both HST and HAP, however ~20% of the HAP solutions have unreliable fit RMS values > 50 mas (see Figure 5).

Gaia eDR3 provides the most accurate alignment compared to other catalogs, with ACS/WFC, WFC3/UVIS, and WFC3/IR aligned to within ~0.3 detector pixels. Of these three most-used detectors, WFC3/UVIS images have the lowest RMS values in milli-arcsec, due to the smaller native plate scale, where the median values are 10 ± 8 mas and 13 ± 8 mas, for HST and HAP products respectively. A small subset of images have RMS values ~300 mas, and the WCS solutions for these observations are likely to be inaccurate and may be reset to use the *a priori* solution. (For details on manual reprocessing, see Section 8). When excluding very large values, the median RMS and the associated dispersion are roughly equal for HST and HAP data products with Gaia eDR3 WCS solutions. One exception is the WFC3/IR detector, where the HAP success fraction is ~10% lower than for HST products, with a median RMS which is 30% larger and a similar dispersion.

For WFC3/UVIS, ACS/WFC, and ACS/SBC observations in Figures 1, 3, & 5 the rolling average **RMS versus date reaches a minimum value between 2016-2018 for observations aligned to Gaia eDR3, which has a reference epoch of J2016.0**. This dip is due to proper motion effects and is more pronounced in HAP products for WFC3/UVIS and ACS/WFC compared to standard products. A similar dip is not seen for either the WFC3/IR detector, which is under-sampled by a factor of ~two, or for the ACS/HRC detector, which ceased operations in January 2007.

5.2 GSC242

While each catalog has a set of images with very large RMS values (suggesting a poor fit), the highest fraction of these occurs for GSC242, and the effect is even more pronounced for HAP data products. For example, the fraction of WFC3/UVIS images with RMS values above 70 mas (see Figure 1) is 25% for HAP compared to 10% for HST products aligned to GSC242. WFC3/IR observations, on the other hand, rarely have large fit RMS values, regardless of which catalog is used for alignment (see Figure 2).

Even after clipping, the median fit RMS for WFC3/UVIS and ACS/WFC is ~2x larger for HAP versus HST images aligned to GSC242 but is relatively the same for WFC3/IR and ACS/HRC. In the plots of fit RMS versus the number of matches for WFC3/UVIS and ACS/WFC, the HAP RMS values have a tighter distribution centered at ~50 mas, whereas the HST values have a large spread of values down to ~10 mas. (A similar effect is not seen for HST products aligned to Gaia eDR3 or 2MASS.)

For HST images obtained in the same visit, the fit RMS can vary widely depending on the signal-to-noise of the observations or the target being observed. Lower RMS values typically correspond to wide filters with many well-measured stars, whereas narrow filters may have fewer stars and larger positional uncertainties due to the extended emission line background. While the large spread in RMS is primarily seen for the ACS/WFC and WFC3/UVIS detectors, which have the largest variety of passbands, it is also noticeable for a subset of WFC3/IR images. The lowest RMS values for GSC242 are similar to the lower envelope of Gaia eDR3 fit residuals, suggesting that this catalog performs exceptionally well for specific targets and filters. Since GSC242 incorporates a variety of ground and space-based catalogs which are tied to DR2 (Lasker et al. 2021), this may also explain the large spread in fit RMS. For example, a subset of the larger fit RMS values may correspond to fields where the less accurate ground-based PanSTARRS sources were matched.

To create the HAP data products, the alignment is based on a combined list of sources from each filter in a given visit, and this may be the reason for the tighter distribution in RMS values for WFC3/UVIS and ACS/WFC, albeit at a higher median value. For filters with few catalog matches, aligning the visit-level observations as a set, rather than individually, may be a better strategy and allow for more uniform detector sampling. In this case, the larger median value of the visit-level RMS (after clipping) is not an indication of poor alignment but instead a reflection of the breadth of sources with varying positional accuracies used for catalog matching.

For ACS/SBC, the median fit RMS of GSC242 is flat at ~20 mas between 2002-2022. Datasets with Gaia eDR3 solutions have similar RMS values as GSC242 near the reference epoch of J2016.0 for eDR3. For the earliest ACS observations in 2002, the rolling average RMS rises slowly to ~30 mas due to proper motion effects. Because this detector has the highest resolution (excluding the inoperable ACS/HRC), the trend in the RMS with date appears more pronounced for ACS/SBC than for other HST detectors.

5.3 2MASS

Only a small fraction of HST *alignable* datasets (<5%) have WCS solutions aligned to 2MASS, and this catalog is typically only used for standard HST products observed with WFC3/UVIS, WFC3/IR, and ACS/WFC. A few dozen HAP products for WFC/UVIS and ACS/WFC are successfully aligned to 2MASS, and these have similar median RMS values as those computed for standard HST products. For ACS/HRC and ACS/SBC, there are currently no 2MASS solutions for either HST or HAP data.

For WFC3/UVIS and ACS/WFC, the rolling average fit RMS values for the purely ground-based 2MASS catalog (red points in Figures 1 and 3) are up to twice as large as those for GSC242. For the highly under-sampled WFC3/IR detector, on the other hand, the fit RMS values for 2MASS WCS solutions are more similar to the RMS for GSC242 solutions. The drop in the WFC3/UVIS rolling average around 2018 (red curve in Figure 1) is likely not significant and a result of fewer datasets with 2MASS solutions at later dates leading to larger uncertainties in the rolling average.

In the case of fewer than ten catalog matches, the scatter in the 2MASS RMS increases rapidly as the number of sources matches decreases. WCS solutions for these datasets should be treated with caution and may need to be refit with a less complex fitting geometry (or reset to the *a priori* WCS) in order to align with additional visits of the same target.

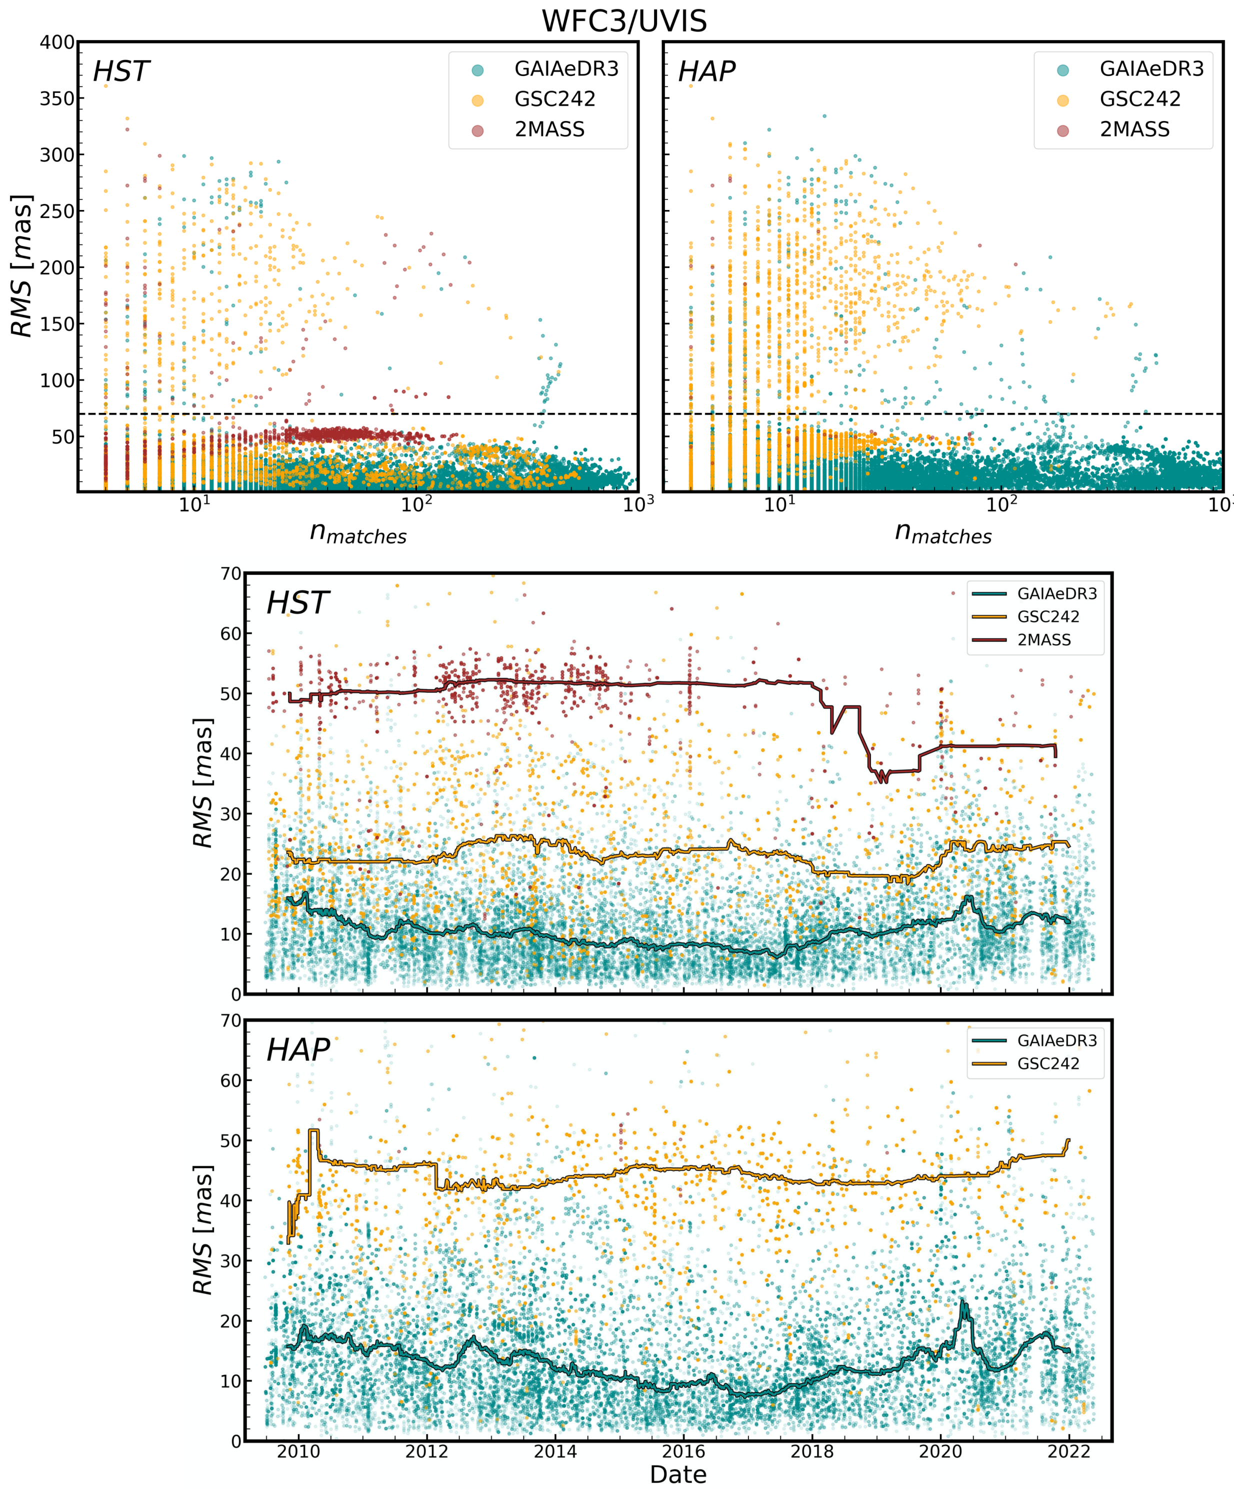

Figure 1: Fit RMS for WFC3/UVIS versus the log of the number of matches to different reference catalogs. The left panel shows the results for standard 'HST' products and the right panel shows advanced 'HAP' products which highlight a subset of images with RMS values > 70 mas. The fit RMS is plotted versus date (2009-2022) for 'HST (center) and 'HAP' (bottom). The solid line is the 'rolling' median in a window of 180 days for eDR3 and 720 days for GSC242 and 2MASS.

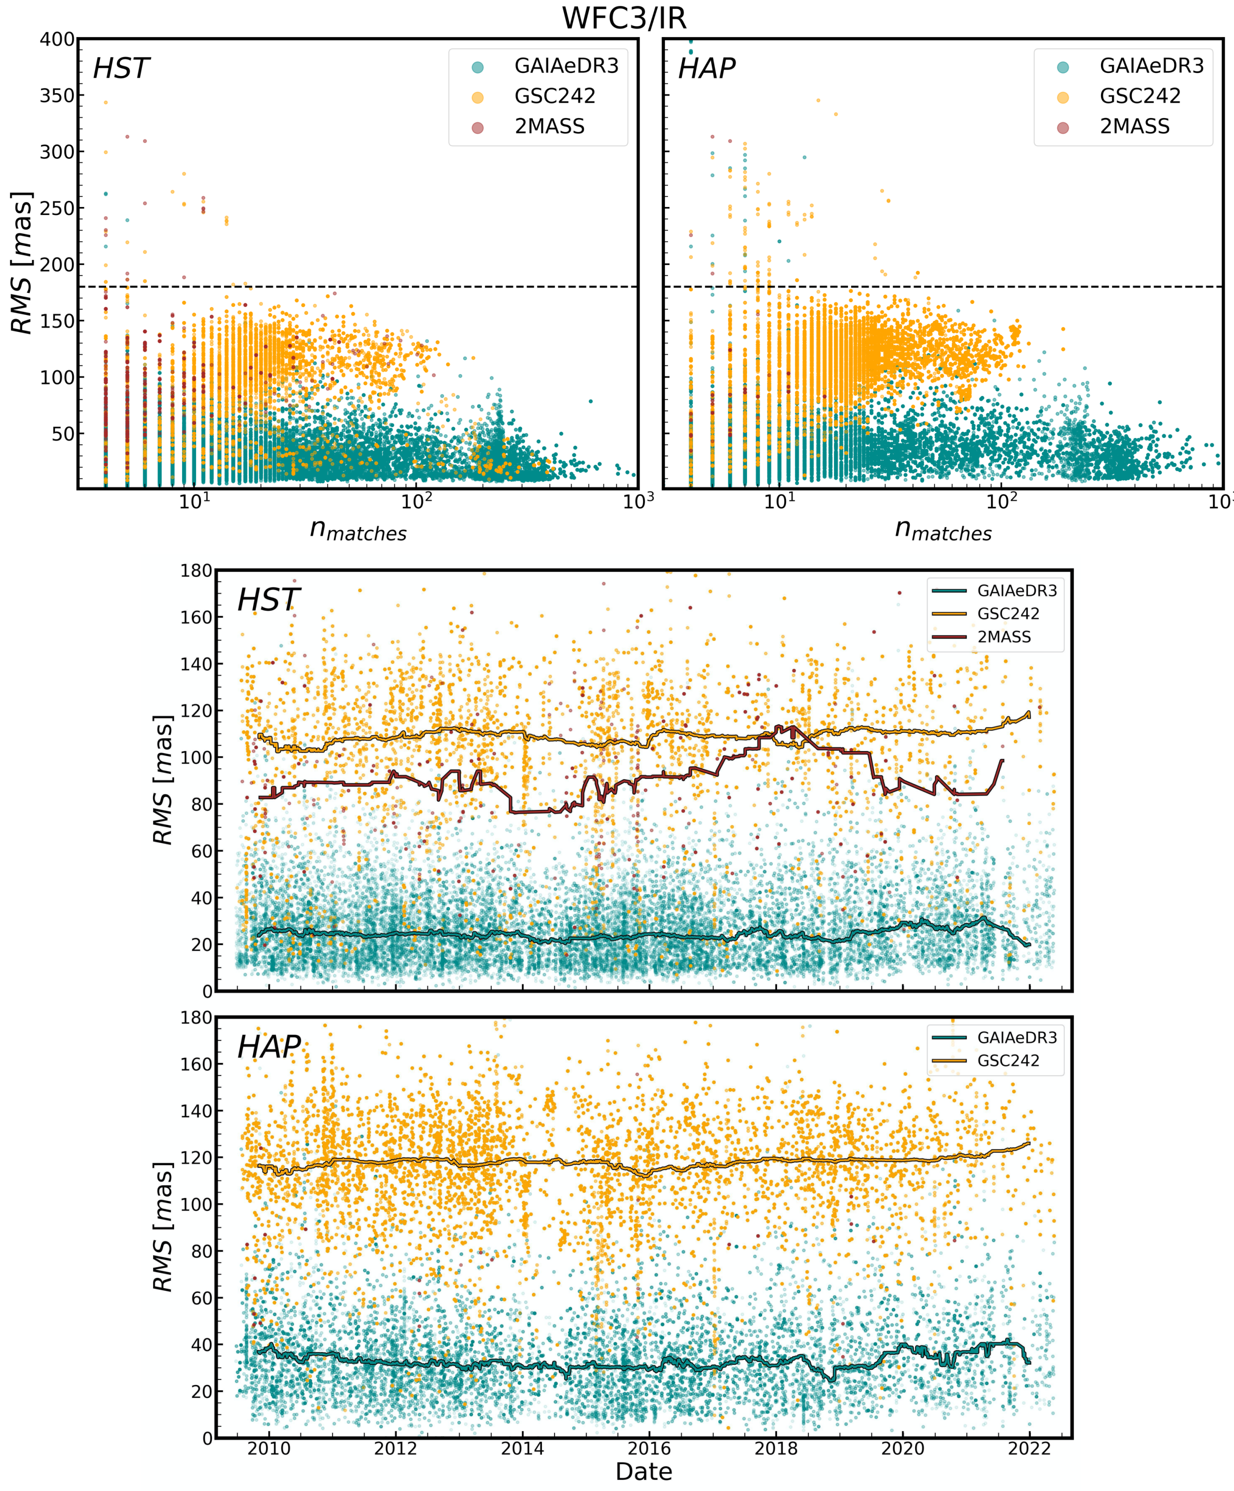

Figure 2: Fit RMS for WFC3/IR versus the log of the number of matches to different reference catalogs. The left panel shows the results for standard 'HST' products and the right panel shows advanced 'HAP' products which highlight a subset of images with RMS values > 180 mas. The fit RMS is plotted versus date (2009-2022) for 'HST' (center) and 'HAP' (bottom). The solid line is the 'rolling' median in a window of 180 days for eDR3 and 720 days for GSC242 and 2MASS.

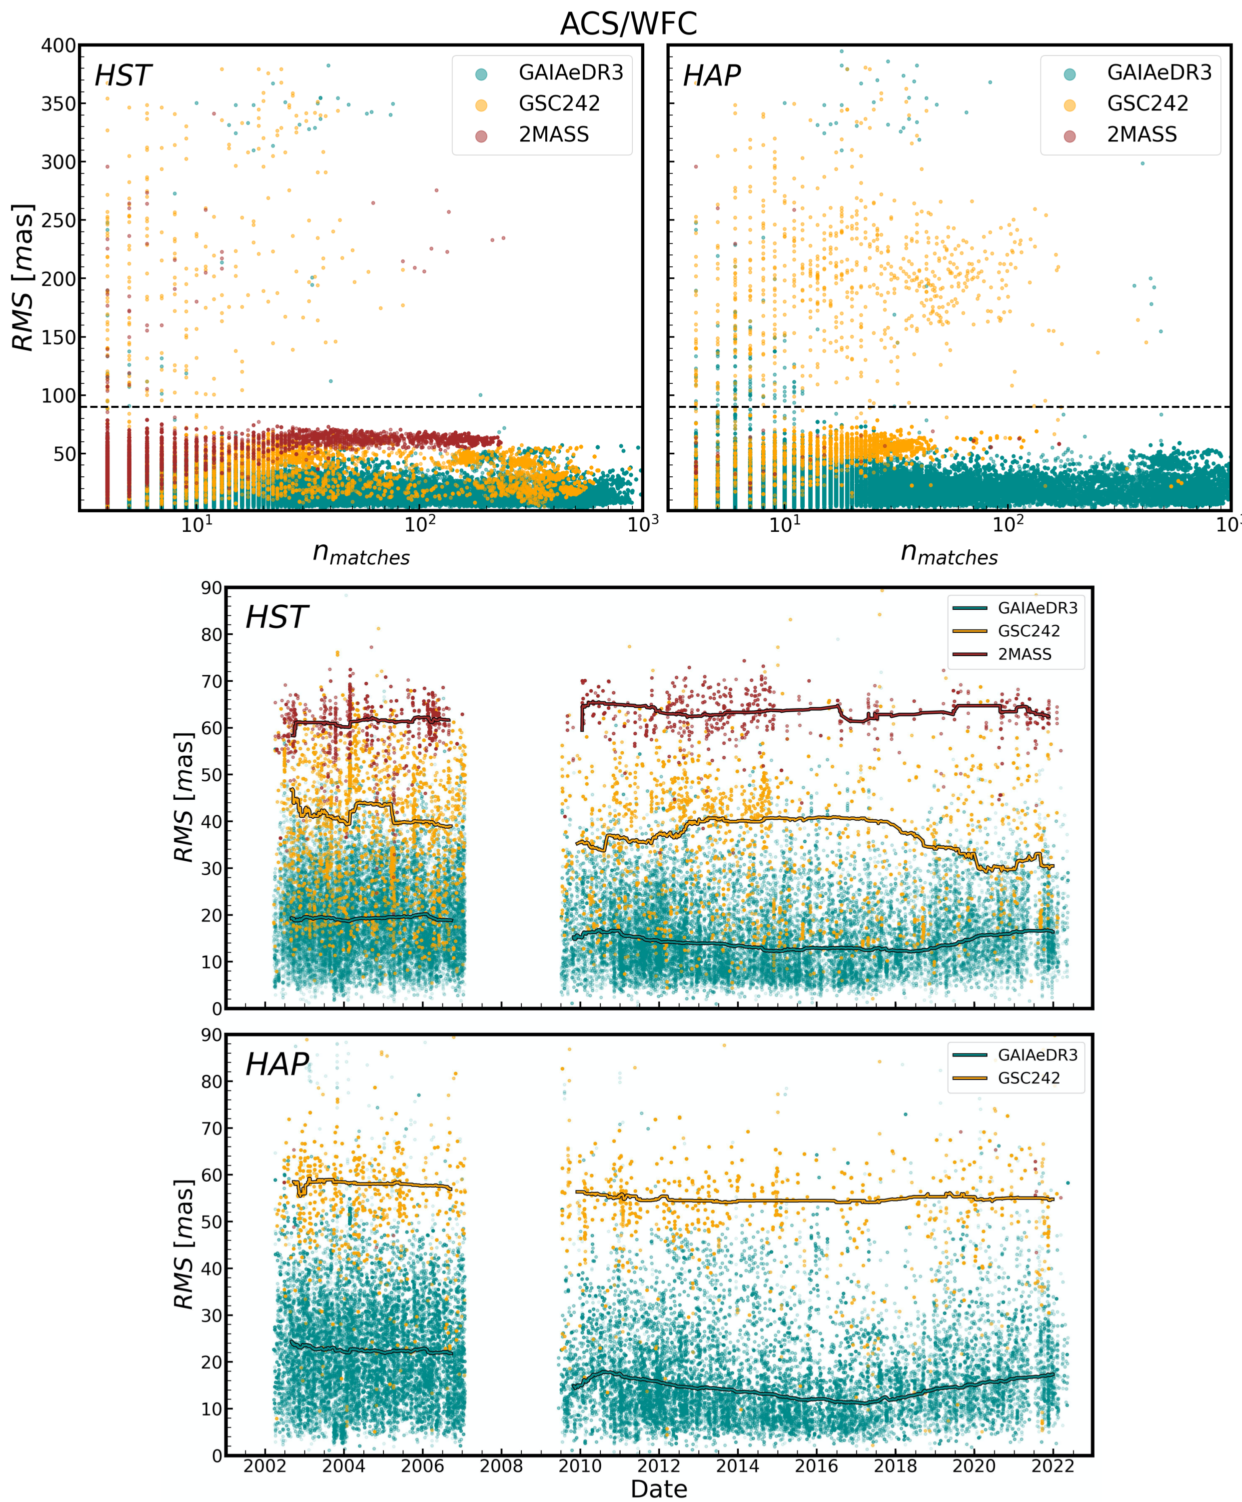

Figure 3: (Top) Fit RMS for ACS/WFC versus the log of the number of matches to different reference catalogs. The left panel shows the results for standard 'HST' products and the right panel shows advanced 'HAP' products highlighting a subset of images RMS values > 90 mas. The fit RMS is plotted versus date (2002-2022) for 'HST' (center) and 'HAP' (bottom). The solid line is the 'rolling' median in a window of 720 days for eDR3 and 1980 days for GSC242 and 2MASS.

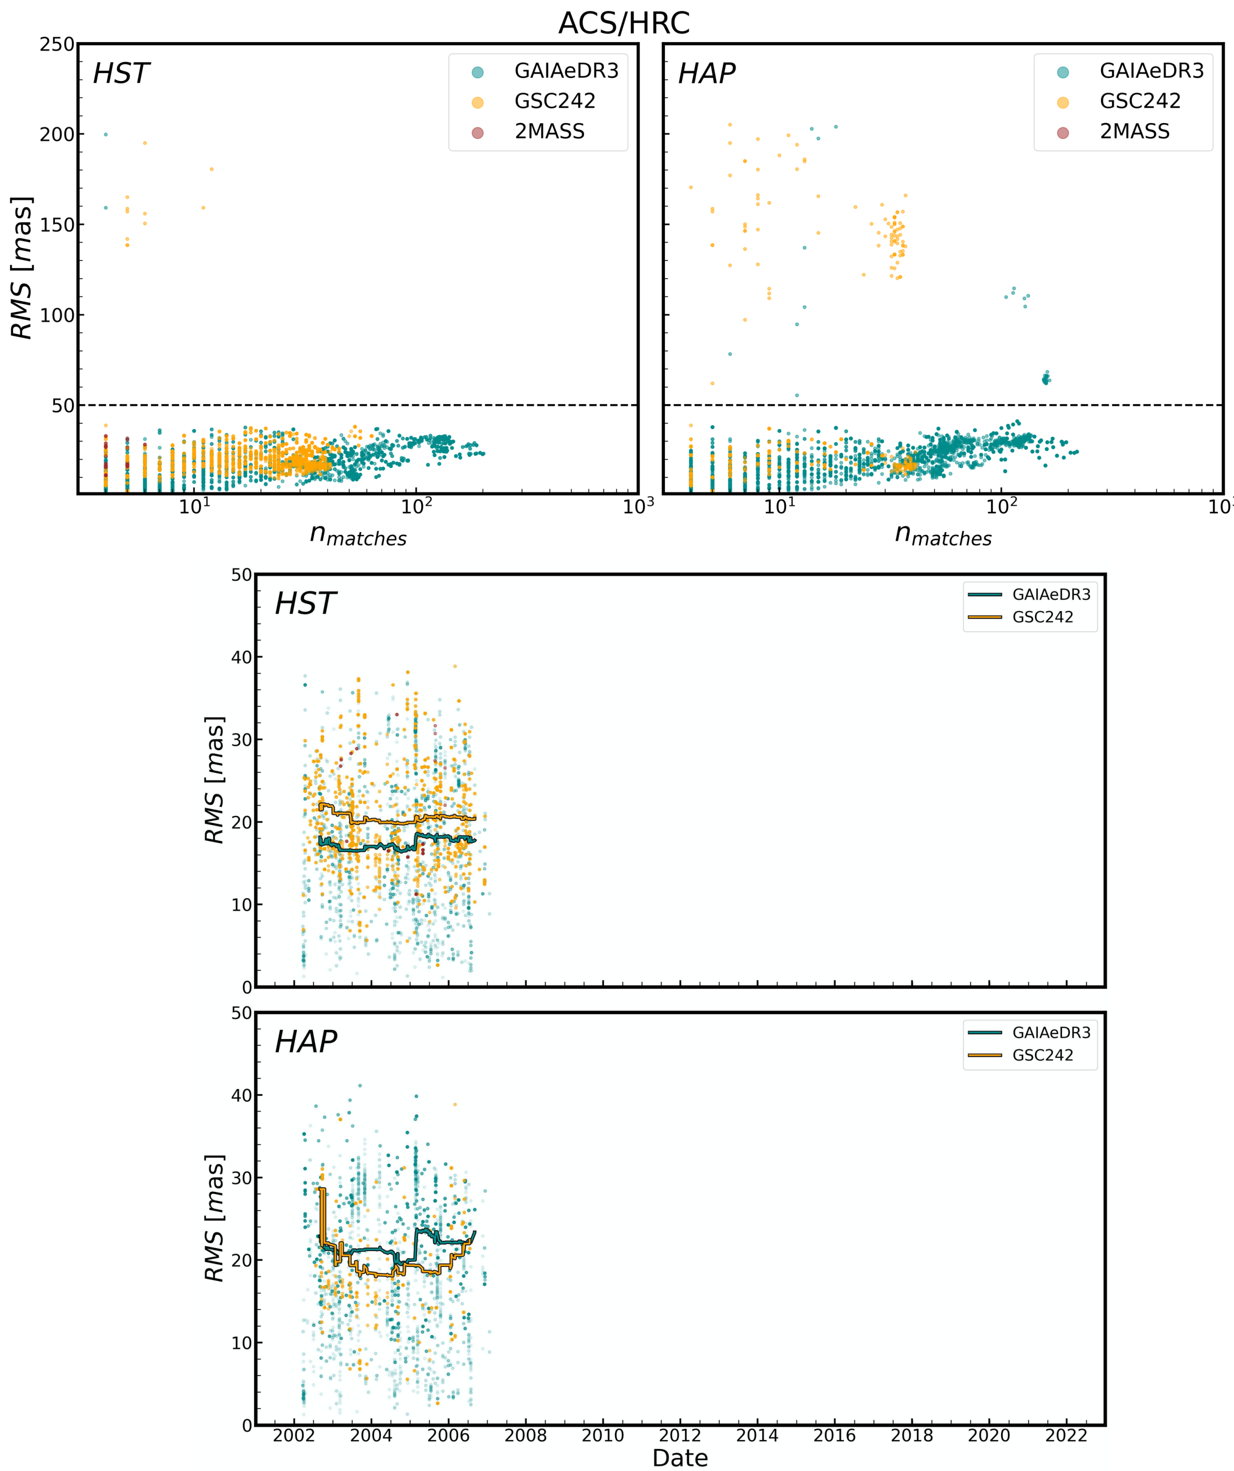

Figure 4: (Top) Fit RMS for ACS/HRC versus the log of the number of matches to different reference catalogs. The left panel shows the results for standard ‘HST’ products and the right panel shows advanced ‘HAP’ products highlighting a subset of images RMS values > 50 mas. The fit RMS is plotted versus date (2002-2022) for ‘HST’ (center) and ‘HAP’ (bottom). The solid line is the ‘rolling’ median in a window of 720 days for eDR3 and 1980 days for GSC242. ACS/HRC ceased operations in January 2007.

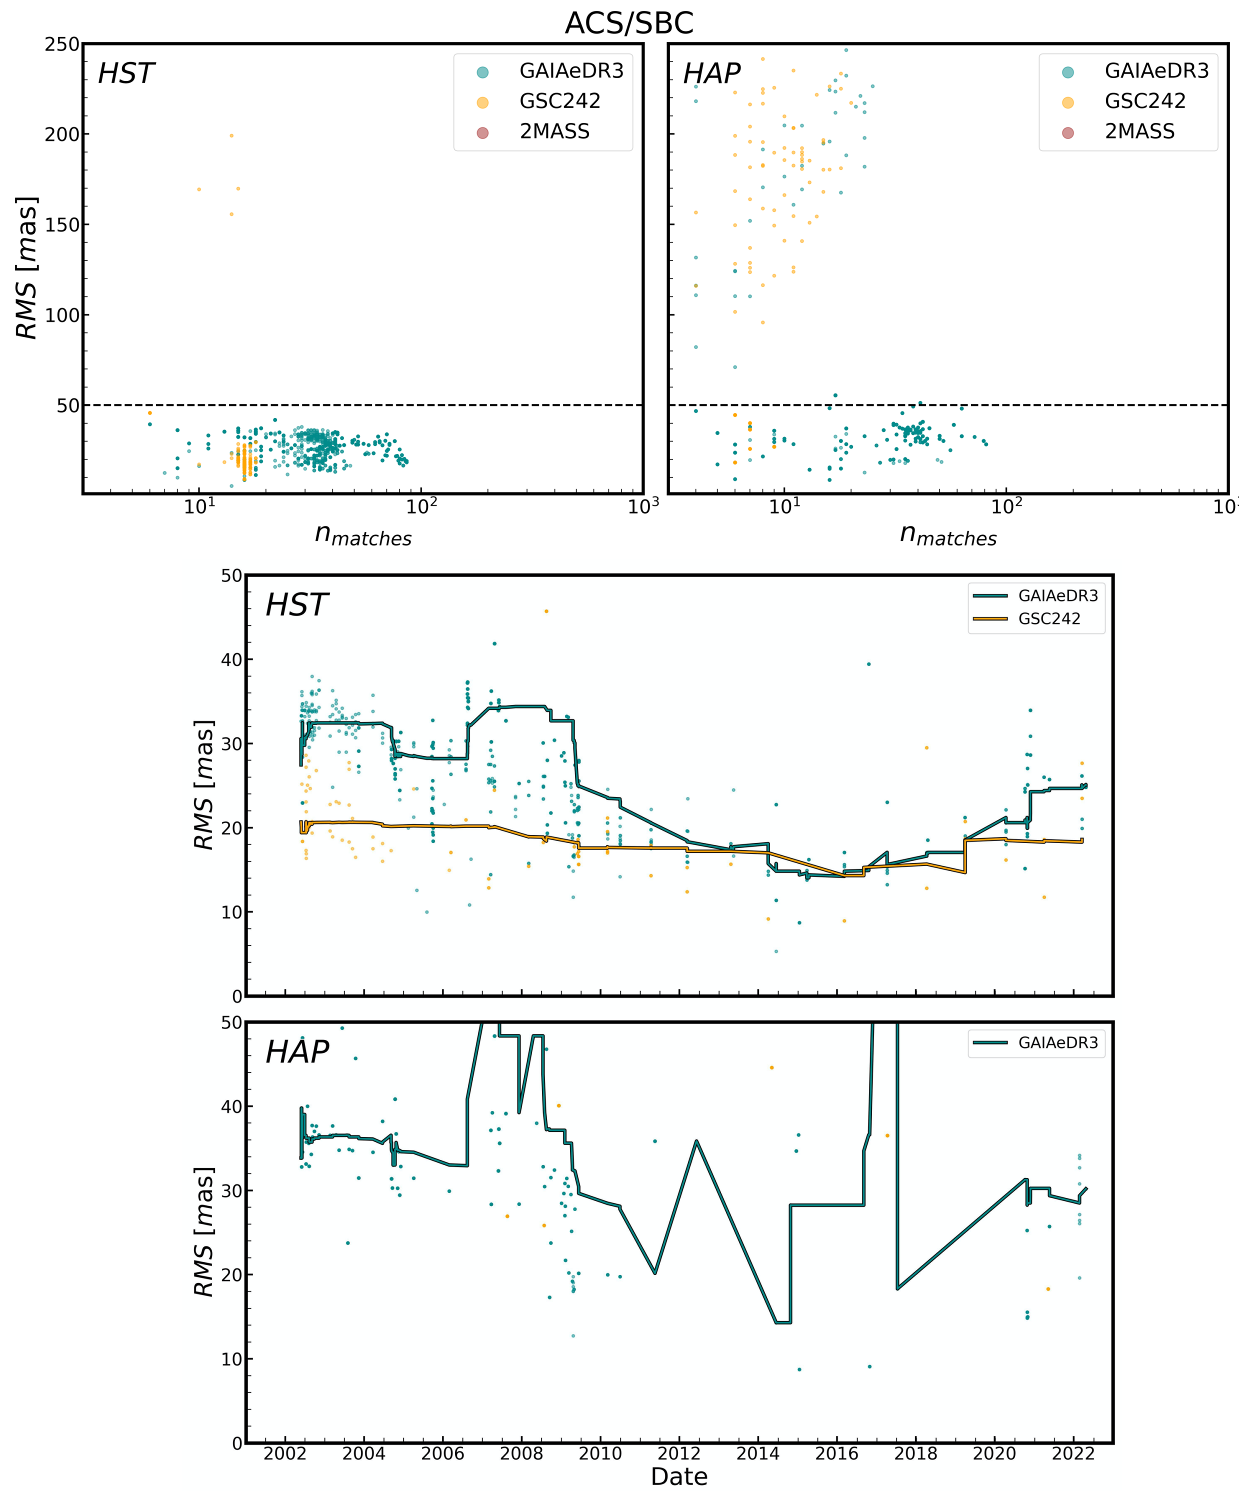

Figure 5: (Top) Fit RMS for ACS/SBC versus the log of the number of matches to different reference catalogs. The left panel shows the results for standard 'Pipeline' products and the right panel shows advanced 'SVM' products highlighting a subset of images RMS values > 50 mas. The fit RMS is plotted versus date (2002-2022) for 'HST' (center) and 'HAP' (bottom). The solid line is the 'rolling' median in a window of 720 days for eDR3 and 1980 days for GSC242.

6. Three ‘Levels’ of MAST Products

6.1 Filenames for HST and HAP (SVM, and MVM) Products

Standard MAST drizzled products have file names which follow the HST naming convention `ipppssoot_drz.fits` , where:

Element	Definition
`i`	Instrument (ACS= ‘j’, WFC3= ‘i’)
`ppp`	Unique identifier for each proposal number
`ss`	Obset ID, typically the two-digit visit ID from the Phase II program
`oo`	Observation number, unique for each exposure
`t`	Transmission source, typically ‘q’ for Solid State recorder (post-SM2)

WFC3/UVIS and ACS/WFC standard data products include additional CTE-corrected drizzled products with the suffix `_drc.fits`. Other products include the calibrated files `_flt.fits`, the calibrated, CTE-corrected files `_flc.fits` files and a corresponding headerlet `_hlet.fits` solution which contains information about the WCS solution and the image distortion correction. Standard products also include the uncalibrated `_raw.fits` files, the image association `_asn.fits`, and the telemetry data `_spt.fits`.

HAP SVM drizzled products have file names which include information about the proposal, e.g. `hst_propid_obsetid_instr_detector_filter_ipppss_drz.fits`, where:

Element	Definition
`proposid`	Proposal ID for the visit
`obsetid`	Two-digit visit ID from the proposal, given above as ‘`ss`’
`instr`	Name of HST instrument from **INSTRUME** header keyword
`detector`	Name of detector from **DETECTOR** header keyword
`filter`	Name of filter from **FILTER** or **FILTER1, FILTER2** header keywords
`ipppss`	First 6 characters of the standard rootname for this visit

For WFC3/UVIS and ACS/WFC, the SVMs use *only* the calibrated, CTE-corrected `_flc.fits` data products as input. Drizzled SVM products `_drc.fits` are produced for each filter, as well as a `_total_drc.fits` image which combines all filters in a given visit per detector. The total image is used a detection image for producing ‘point’ and ‘segment’ catalogs (see Section 7). ‘Exposure level’ products from single frames are also produced for each visit, and these have similar names as the ‘filter level’ and ‘total level’ products, except for having two extra characters in the rootname: `ipppssoo` (see Table 5).

HAP MVM drizzled products have a file names which reflect the sky cell naming convention:
`hst_skycell-pPPPPxXXyYY_instr_detector_filter_label_drz.fits`

Element	Definition
`PPPP`	**Projection Cell** ID as a zero-padded 4-digit integer
`XX`, `YY`	**Sky Cell** ID within Projection Cell as zero-padded 2-digit integers
`instr`	Name of HST instrument from **INSTRUME** header keyword
`detector`	Name of detector from **DETECTOR** header keyword
`filter`	Name of filter from **FILTER** (or FILTER1, FILTER2) header keywords
`label`	Layer-specific designation describing the output plate scale: **'all'** for layers at **0.04″ scale,** or **'coarse-all'** for WFC3/IR at **0.12″ scale**

MVM products include only the drizzled FITS files for each detector, referred to as 'layers', and their corresponding log files. WFC3/IR layers are created at two scale values: a 'coarse' grid of 0.12″/pix, similar to the native plate scale, and a fine grid of 0.04″/pixel (super-sampled x3) to match the native WFC3/UVIS plate scale. For ACS/WFC images, the 0.04″/pixel grid is slightly smaller than the native 0.05″/pixel detector scale. Future processing may enable the creation of additional layers based on criteria such as exposure time, observation date, or catalog solution.

Due to the difficulty of aligning UV sources to the selected set of reference catalogs, MVMs are not produced for ACS/HRC or ACS/SBC. These detectors have a very small field of view and are typically not used for mosaicking. Additionally, the native scale of these detectors is smaller than the 0.04″/pixel resolution produced for MVMs, so we advise users interested in HRC or SBC observations to use either the standard HST or SVM data products, depending on the reference catalog and alignment quality (see Section 5). Alternately, users may use the STScI-provided tools to create custom MVM layers for these detectors (see Section 8).

An example of the 'HST' and 'HAP' data products corresponding to a single visit of WFC3/IR observations of EMACS J0241.0+2557 is provided in Table 5, while a more complex example for the same target is provided in Appendix A. This galaxy cluster lies within sky cell `p1857x08y11` and was observed in multiple HST proposals, detectors, and bandpasses. Table A1 lists the MVM drizzled data products for the RGB color image in Figure A1. The separate layers are assigned different colors for ACS/WFC F606W (blue), WFC3/IR F110W (green), and WFC3/IR F140W (red). The exposures are aligned to Gaia eDR3 and drizzled to the same sky cell footprint with a pixel scale of 0.04″/pixel for easy comparison.

Table A2 lists the SVM data products for the same sky cell, and these include both the calibrated and drizzled images. The exposure-level, calibrated images `_flt.fits` (WFC3/IR) and the CTE-corrected images `_flc.fits` (ACS/WFC) are used as input to create the MVMs listed in Table A1, combining all visits in a given filter, regardless of date or WCSNAME solution. SVM products include 'point' and 'segment' photometry for each filter and for the detection image, named `_total*_ dr[cz].fits`, computed for separately for ACS and WFC3. Table A3 lists the standard HST products associated with this sky cell which include the drizzled images for each filter (with and without CTE-correction), the exposure-level calibrated images, `_fl[tc].fits`,

Table 5: Three 'levels' of MAST data products for WFC3/IR Visit 01 from program 15132. Non-indented rows in the first column show 'minimum recommended products', while indented rows show additional available FITS files. The HAP SVM products include drizzled images plus both point and segment catalogs. The HAP MVM products include a drizzled 'layers' at a scale of 0.04″/pix and 0.12″/pix, similar to the native detector scale.

MAST FITS Data Products	Description	Rootname
idkl0l020_drz.fits	HST DRZ, 0.128″/pix	idkl0l020
idkl0l020_asn.fits	Association	idkl0l020
idkl0ldmq_flt.fits	FLT Exp 1	idkl0ldmq
idkl0ldmq_ima.fits	IMA Exp 1	idkl0ldmq
idkl0ldmq_raw.fits	RAW Exp 1	idkl0ldmq
idkl0ldoqq_flt.fits	FLT Exp 2	idkl0ldoq
idkl0ldoqq_ima.fits	IMA Exp 2	idkl0ldoq
idkl0ldoqq_raw.fits	RAW Exp 2	idkl0ldoq
hst_15132_0l_wfc3_ir_f140w_idkl0l_drz.fits	SVM DRZ,0.128″/pix	idkl0l
hst_15132_0l_wfc3_ir_f140w_idkl0l_point-cat.ecsv	Point Source Catalog	idkl0l
hst_15132_0l_wfc3_ir_f140w_idkl0l_segment-cat.ecsv	Segment Catalog	idkl0l
hst_15132_0l_wfc3_ir_f140w_idkl0ldm_flt.fits	FLT Exp 1	idkl0ldm
hst_15132_0l_wfc3_ir_f140w_idkl0ldm_drz.fits	DRZ Exp 1	idkl0ldm
hst_15132_0l_wfc3_ir_f140w_idkl0ldo_flt.fits	FLT Exp 2	idkl0ldo
hst_15132_0l_wfc3_ir_f140w_idkl0ldo_drz.fits	DRZ Exp 2	idkl0ldo
hst_skycell-p1857x08y11_wfc3_ir_f140w_all_drz.fits	MVM DRZ, 0.04″/pix	p1857x08y11
hst_skycell-p1857x08y11_wfc3_ir_f140w_coarse-all_drz.fits	0.12″/pix	p1857x08y11

the 'intermediate' WFC3/IR products `_ima.fits` (prior to ramp fitting), and the uncalibrated raw images `_raw.fits`. The calibrated HST standard pipeline images are used as input to create the SVM drizzled products listed in Table A2, after an improved relative alignment across filters in a given visit, followed by a global catalog realignment.

HAP data are reprocessed when the underlying HST data products are reprocessed, so they are always archived with the most recent calibration software and reference files. Additionally, **the tools for creating advanced products are included as part of the `DrizzlePac` software distribution so that users can create their own custom mosaics** (see Section 8).

6.2 Drizzle Parameters

Custom parameters for `AstroDrizzle` have been optimized for each HST detector and were selected based on the number of input frames and the filters of interest. For standard drizzled products, the parameter settings are pulled from the `MDRIZTAB` reference file, e.g. see Mack & Bajaj (2019) for WFC3/IR and Hoffmann & Avila (2017) for ACS/WFC. In particular, the cosmic ray rejection parameters `driz_cr_snr` and `driz_cr_scale` are fine-tuned for each detector in order to avoid flagging the center of point sources in the case of small ~0.1 pixel alignment errors. In addition, the data quality (DQ) flags are defined differently for each HST detector, so the recommended parameters for including or ignoring specific bit flag values (`driz_sep_bits, final_bits`) will also differ. To reproduce the settings used in the pipeline when reprocessing observations, the `AstroDrizzle` parameter `MDRIZTAB` may be set to 'True' so that the recommended parameters will be used. Additional modifications to the parameter settings may then be entered via the command line (see Section 8).

The drizzle parameters for advanced products (SVMs and MVMs) may be found in the `DrizzlePac` repository 'hap_pars'[7]. In particular, the sky background is not removed from the drizzled images, but instead 'equalized' across all exposures by computing differences in the median value for overlapping pixels. Standard drizzled products for WFC3/UVIS and WFC3/IR were updated in 2019 to use this strategy (e.g. `skymethod='match'`) via a new `MDRIZTAB` reference file. The ACS/WFC and ACS/HRC reference file, however, still uses the older '`localmin`' setting which separately measures and removes the sky for each exposure. For large, extended targets which fill most of the field of view, this method tends to overestimate the sky background.

For MVM data, the drizzle parameters are set so that only the 'sky matching' and 'final drizzling' steps in `AstroDrizzle` are performed. (Cosmic ray flags, DQ = 4096, are adopted directly from the SVM calibrated images, so all other steps are turned off.) In order to avoid resetting the 4096 DQ flags in the input data, the parameter `resetbits` has been set to a value of `0`.

6.3 Header Keywords

After reprocessing, new keywords are populated in the FITS image headers, and these may be used to assess the quality of the alignment or to determine whether different datasets share the same WCS solution. When *a posteriori* WCS solutions are computed, the header contains statistics on the number of matched sources used in the fit, the fitting geometry used, and the RMS of the fit (see Table 7). Because the HAP drizzled images may contain large, empty regions with `nan` values, new keywords are added to the `_dr[zc].fits` science header. These include the fraction of pixels with data, the mean & median number of input frames, the mean & median exposure time per pixel, and the mean & median value in the drizzled weight (WHT) image.

While the header keyword `EXPTIME` reports the total exposure time of all input frames, regardless of overlap, this is a less useful quantity for advanced products. When combining frames, the software weights each image by individual exposure time, so the weight provides an estimate of the exposure depth in seconds. Thus the `MEANEXPT` is good for knowing the mean exposure time for input pixels (FLT) with data, while the `MEANWHT` is useful for determining the mean exposure for output (DRZ) pixels with data. Note that MVM layers for WFC3/IR are produced for two output scale values: 0.12″/pixel and 0.04″/pixel. While the former 'coarse' image will have a `MEANWHT` value similar to the total exposure time, the more finely-sampled image will have a `MEANWHT` value which is smaller by a factor of 9.

[7] https://github.com/spacetelescope/drizzlepac/tree/master/drizzlepac/pars/hap_pars

Table 6: Header keywords describing the WCS solutions used for alignment.

Keyword	Definition	HST flt/flc	HST drz/drc	HAP flt/flc	HAP drz/drc
WCSNAME	Coordinate System Title	x	x	x	x
WCSTYPE	Coordinate System Description	x	x	x	x
HDRNAME	Label for headerlet solution	x	x	x	x

Table 7: Header keywords related to catalog alignment for observations with '*a posteriori'* WCS solutions which are 'FIT' to a reference catalog.

Keyword	Definition	HST flt/flc	HST drz/drc	HAP flt/flc	HAP drz/drc
NMATCHES	Number of matched sources in fit	x	x	x	x
FITGEOM	Geometry used for fitting	x	x	x	x
RMS_RA	RMS in RA for fit (milli-arcsec)	x	x	x	x
RMS_Dec	RMS in Dec for fit (milli-arcsec)	x	x	x	x
CRDER1	RMS in RA for fit (degrees)	x	x	x	x
CRDER2	RMS in Dec for fit (degrees)	x	x	x	x

Table 8: Header keywords for 'HAP' drizzled data, including a new classification 'level'. Additional keywords provide the mean (median) exposure time and number exposures for all pixels with valid input data, as well as the mean (median) value in the output drizzled weight image which provides the total exposure time per pixel.

Keyword	Definition	HST flt/flc	HST drz/drc	HAP flt/flc	HAP drz/drc
HAPLEVEL	Classification 'level' (1=single, 2=combined, 3=mosaic)				x
SCELLID	ID of the SkyCell				x
NPIXFRAC	Fraction of pixels with data				x
MEANEXPT	Mean exposure time per pixel with data				x
MEDEXPT	Median exposure time per pixel with data				x
MEANNEXP	Mean # exposures per pixel with data				x
MEDNEXP	Median # exposures per pixel with data				x
MEANWHT	Mean WHT value per pixel with data				x
MEDWHT	Median WHT value per pixel with data				x

7. Source Catalogs

SVM data products include two types of source catalogs, referred to as 'point source' and 'segment' catalogs. Both are generated using utilities from `Photutils`[8] (Bradley et al. 2022) with the point source catalog created using functionality similar to `DAOPhot` (Stetson 1987), and the segment catalog created using functionality similar to `Source Extractor` (Bertin & Arnouts 1996). These new catalogs are intended to replace the similar HLA source lists for ACS and WFC3. (The HLA source lists will continue to be available but include only HST observations acquired through October 2017.) Note that the HLA products also include WFPC2 images and catalogs, which are not yet available as HAP data products but which are planned for a future data release. A new version of the Hubble Source Catalog[9] is currently in development and will be built from a combination of the HAP source lists and the HLA WFPC2 source lists.

The detection thresholds used to generate point and segment catalogs were fine-tuned to roughly match the number of objects in the HLA catalogs, where the latter used a 3-sigma detection in most cases. (For convenience, the precise threshold values for each detector are written in the text header of each HAP catalog.) Figure 6 compares the number of sources in each HAP catalog (y-axis) to the number of sources in the corresponding HLA catalog (x-axis) for a test sample of ~48,000 datasets, where the left panel is the point catalog and the right is the segment catalog. While the WC3/IR segment catalogs are slightly shallower than the HLA, the general depth of the catalogs is similar. Figure 7 compares the fraction of sources matching between HLA and HAP catalogs as a function of the minimum number of sources in either HLA or HAP and confirms that there are no obvious source quality issues (e.g., spurious detections from artifacts).

The software used to produce source catalogs is provided on the `DrizzlePac` 'Catalog Generation' page[10], along with a description of their contents. Photometry is computed for each HST detector in two small apertures, `Aper1` and `Aper2`, which are listed in Table 9. Photometry is provided in the ABMAG system and calibrated using keywords populated in the image headers which corresponding to an 'infinite' aperture enclosing all of the light. **To convert to total magnitudes, aperture corrections must be applied to account for flux falling outside of the selected aperture**.

To compute aperture corrections, a two-step process is recommended. **The first corrects photometry from some small aperture to a larger 'standard' aperture for each detector**, the radius beyond which changes in the telescope focus, orbital breathing effects, or spatial variations in the PSF do not affect the aperture photometry. When possible, this correction should be measured directly from isolated stars in the drizzled science frames.

The second step corrects the total brightness from the 'standard' aperture to an 'infinite' aperture, which encloses all of the light and is may be taken from encircled energy (EE) tables. These tables are provided by the instrument teams for each detector and are derived from high signal-to-noise observations of isolated stars out to large radii. EE tables for the different

---

[8] https://photutils.readthedocs.io
[9] https://archive.stsci.edu/hst/hsc
[10] https://drizzlepac.readthedocs.io/en/latest/catalog_generation.html

detectors may be found on the 'Data Analysis' pages for ACS[11] and WFC3[12]. For further detail on aperture corrections, see Section 5.1.2 of the ACS Data Handbook[13] (Lucas & Ryan et al. 2022) and Section 9.1.8 of the WFC3 Data Handbook[14] (Sahu et. al. 2021).

Columns 7 and 8 of Table 9 list the 'standard' and 'infinite' apertures for each detector. For ACS/WFC and ACS/HRC, for example, aperture photometry in a radius less than the standard 0.5" aperture will vary with time and detector position, and the aperture correction for either the HSC or HAP catalog magnitudes requires the two steps described above. ACS/WFC catalog magnitudes corresponding to `Aper2` require a correction from 0.15" (3 pixels) to 0.5" (10 pixels) which may be empirically determined from the drizzled data. Next, the correction from 0.5" to infinity (5.5") may be taken directly from the ACS encircled energy tables for the filter of interest.

For the WFC3/UVIS and WFC3/IR detectors, aperture photometry in a radius less than 0.4" will also vary both with time and detector position. For example, WFC3/UVIS catalog magnitudes corresponding to `Aper1` require a correction from 0.05" (1.3 pixels) to the standard aperture at 0.4" (10 pixels) and then a second correction to 'infinite' aperture (defined at 6") which may be taken from the EE tables for a given filter and CCD (chip).

While there is no formally documented 'standard' aperture for ACS/SBC, the absolute flux calibration was derived using photometry in an aperture radius of 0.5", which corresponds to 20 pixels for observations drizzled to a scale of 0.025"/pixel (Avila et al. 2019). The photometry was then corrected to an infinite aperture of 4.0" using the encircled energy corrections.

Table 9: Summary of SVM data products, including the drizzled scale for each detector (column 2) and the two apertures used to generate source catalogs in arcseconds (columns 3-4) and in drizzled detector pixels (columns 5-6). Aperture corrections are required to obtain total magnitudes, and the 'standard' aperture and 'infinite' aperture used for the absolute photometric calibration are listed for each detector in columns 7 and 8 (see text for details).

HST Detector	Drizzled Scale	HSC Aper1	HSC Aper2	HSC Aper1	HSC Aper2	'Standard' Aperture	'Infinite' Aperture
WFC3/IR	0.128"/pix	0.15"	0.45"	1.2 pix	3.5 pix	0.40"	6.0"
WFC3/UVIS	0.040"/pix	0.05"	0.15"	1.3 pix	3.8 pix	0.40"	6.0"
ACS/WFC	0.050"/pix	0.05"	0.15"	1.0 pix	3.0 pix	0.50"	5.5"
ACS/HRC	0.025"/pix	0.03"	0.125"	1.2 pix	5.0 pix	0.50"	5.5"
ACS/SBC	0.025"/pix	0.07"	0.125"	2.8 pix	5.0 pix	0.50"	4.0"

[11] https://www.stsci.edu/hst/instrumentation/acs/data-analysis/aperture-corrections
[12] https://www.stsci.edu/hst/instrumentation/wfc3/data-analysis/photometric-calibration
[13] https://hst-docs.stsci.edu/wfc3dhb
[14] https://hst-docs.stsci.edu/acsdhb

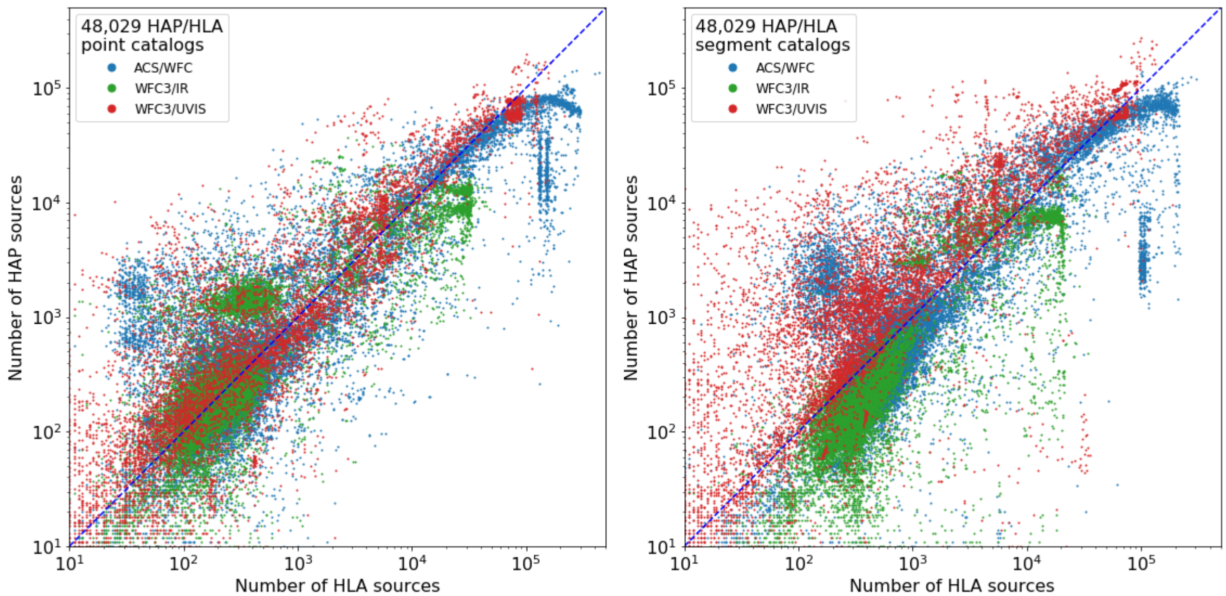

Figure 6: Comparison of number of sources in the HLA and HAP point catalogs (left) and segment catalogs (right) for 48,029 catalogs that appear in both collections. Similar plots were used to fine-tune the HAP detection threshold values for each detector to roughly match with the existing HSC source lists.

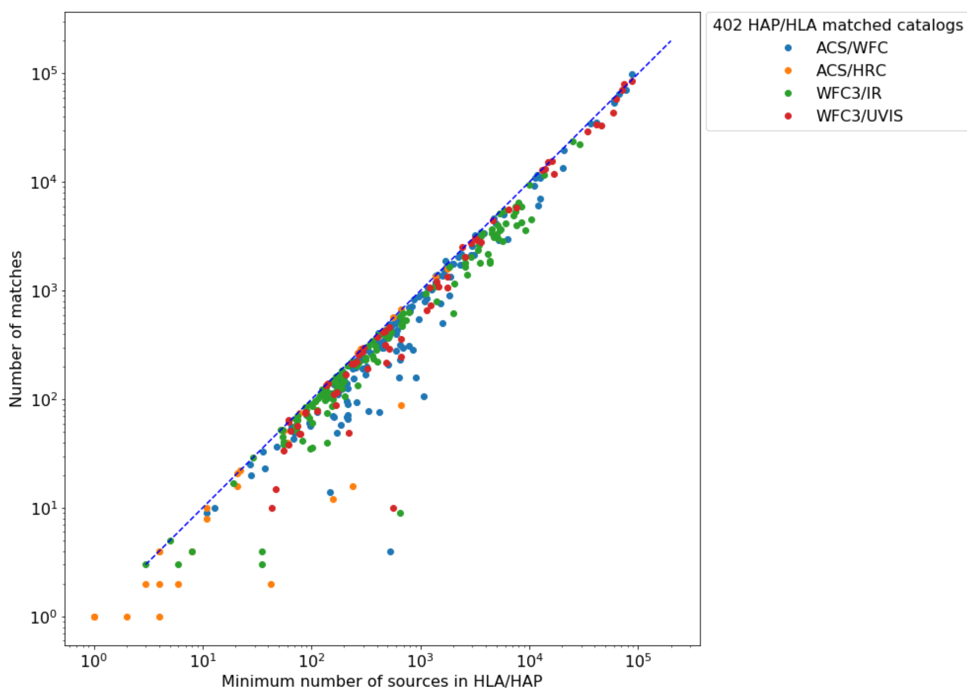

Figure 7: Number of cross-matches between the HLA and HAP, including both point and segment catalogs, versus the minimum number of sources in either HLA or HAP. (A perfect match would put all the the points along the diagonal line.) For a test sample of 402 catalogs, most of the results are close to the blue line confirming that there are no obvious source quality issues (e.g., spurious detections from artifacts).

8. Reprocessing Tools

Observations downloaded from the MAST archive after reprocessing will have extra extensions or 'headerlets' added to the FITS files. These represent different WCS solutions which were derived as the data were (re-)processed, with the 'best' solution applied as the active or primary WCS. Alternately, any of the WCS solutions may be downloaded as separate headerlet files from MAST and then applied to datasets which were downloaded at an earlier date.

A Jupyter notebook titled '*Using updated astrometry solutions*' was developed to familiarize users with the structure of the FITS images and to demonstrate how the primary WCS may be changed to any other preferred solution by applying the headerlet files. This tutorial is available from the `DrizzlePac` notebooks repository[15], which contains additional examples on alignment and drizzling of HST observations.

For datasets with poor catalog alignment, users may choose to reset the WCS to the *a priori* solution with the `stwcs` sub-package `updatewcs`[16]. When the parameter `use_db` is set to `False`, the software will *not* adopt the primary WCS defined in the MAST astrometry database. The WFC3 example below shows how to reset the WCS to the *a priori* solution. The distortion reference files are assumed to be in a directory called '`iref`', which must be set to a local path containing those files.

```
>> import os
>> os.environ['iref']= 'local_path_to_reffiles'
>> from stwcs import updatewcs
>> updatewcs.updatewcs('*flt.fits', use_db=False)
```

Software to improve the image alignment and/or verify the alignment accuracy may be found on the `DrizzlePac` 'Astrometry API' page[17]. This software is able to automatically connect to the HST astrometry database to retrieve and apply headerlets. Python functions for creating, updating, and applying headerlets to FITS images are described via the 'Headerlet User Interface' page[18].

Alternately the relative image alignment may be improved using `TweakReg`[19] and then re-aligned to a catalog. A helpful tutorial titled *'Aligning HST Images to an Absolute Reference Catalog'*[20] is available in the `DrizzlePac` notebook repository. This notebook is based on a WFC3 report by Bajaj (2017) and contains a simplifed version of the original notebook in the *'Gaia Alignment'* repository[21]. Together, these two notebooks show how to query select stars within the HST footprint from the Gaia catalog with positional errors smaller than a user-defined threshold. The catalog may then be used as input to `TweakReg` in order to improve the absolute astrometry prior

[15] https://github.com/spacetelescope/notebooks/tree/master/notebooks/DrizzlePac
[16] https://stwcs.readthedocs.io/en/latest/updatewcs.html
[17] https://drizzlepac.readthedocs.io/en/latest/astrometry_api.html
[18] https://stwcs.readthedocs.io/en/latest/headerlet_ui.html
[19] https://drizzlepac.readthedocs.io/en/latest/tweakreg.html
[20] https://github.com/spacetelescope/notebooks/tree/master/notebooks/DrizzlePac/align_to_catalogs
[21] https://github.com/spacetelescope/gaia_alignment/

to drizzling, with the option of over-plotting the selected Gaia sources over the drizzled image to check the alignment.

Once the alignment has been fine-tuned and the WCS updated, the calibrated observations may be re-drizzled using the recommended `AstroDrizzle` parameters for each detector, as discussed in Section 6.2. The software for reproducing standard HST pipeline drizzled products is described on the 'Running AstroDrizzle' page[22]. To re-drizzle standard HST data products with custom parameters, users are advised to start with the detector-specific settings in the `MDRIZTAB` reference file, which contains unique values based on the number of input frames, filters, etc. and which is available for download from MAST. Any modifications to these parameter settings may then be entered via the command line. In the following example, the software will use the set of calibrated `_flt.fits` images in the user's local directory to determine the recommended parameter values, and the user may then provide any additional non-default parameters values, e.g. drizzling with North oriented up.

```
>> from drizzlepac import astrodrizzle
>> astrodrizzle.AstroDrizzle('*flt.fits', mdriztab=True,\
       final_wcs=True, final_rot=0)
```

Drizzle parameters for advanced products (SVMs and MVMs) are not captured in any reference file, but are instead provided in the `HAP_PARS`[23] section of the `DrizzlePac` repository. The software API for reproducing Hubble Advanced products in MAST is described on the pages for for SVM[24] or MVM[25] data. The software API for generating custom MVM mosaics, e.g. reprocessing data from multiple visits that span multiple sky cells is described on the following 'Make Custom Mosaic' page[26].

Additional utilities are available for 'Manipulating MVM Products'[27]. The `hapcut_utils.py` module contains utilities which may be imported into other Python programs or a Python session for interactive data analysis. The utilities are wrappers around lower-level functionality which perform the core tasks. The core tasks referenced here are `astroquery`, used for MAST inquiries and retrievals, and `astrocut`, used to create the cutouts and cutout combinations.

Note that MAST provides a *separate* API for requesting custom MVM cutouts of drizzled images or for combining images in the same projection cell that span multiple sky cells. For more information about these tools, see the 'HAPcut.MAST API' webpage[28]. An HLA-style interface is currently being developed to aid users in visualizing the MVM layers for various detectors and passbands within the same footprint on the sky. These will be similar to the HLA 'Level 4 Data ' products which allow users to create RGB images by selecting (or deselecting) the observed filters of interest, with the option of overplotting the Gaia, 2MASS, or GSC catalogs, and/or the SVM

[22] https://drizzlepac.readthedocs.io/en/latest/runastrodriz.html
[23] https://github.com/spacetelescope/drizzlepac/tree/master/drizzlepac/pars/hap_pars
[24] https://drizzlepac.readthedocs.io/en/latest/runsinglehap.html
[25] https://drizzlepac.readthedocs.io/en/latest/multivisit_api.html
[26] https://drizzlepac.readthedocs.io/en/latest/makecustommosaic.html
[27] https://drizzlepac.readthedocs.io/en/latest/mvmutilities.html
[28] https://mast.stsci.edu/hapcut/

point and segment catalogs. A test version of these interactive visualization tools was used to create the MVM color-composite images shown in the Appendix.

9. Summary

MAST data products have been updated include improved absolute astrometry in the image header World Coordinate System (WCS) for ACS and WFC3. These include updates to the coordinates of the guide stars used at the time of observation to match GSC v2.40, referred to as an *a priori* WCS. Additionally, observations for each filter may be individually aligned to a select set of reference catalogs and given an *a posteriori* WCS. Standard HST data products follow the familiar naming convention '*ipppssoot_drz.fits*'. If desired, the WCS for these datasets may be easily changed to a different solution or reset to use the *a priori* solution.

Two new types of Hubble Advanced Products (HAP) are also available, and these have improved relative alignment across filters by cross-matching sources in the HST images, followed by a global catalog alignment. The first type of products are referred to as Single Visit Mosaics '*hst_prop_visit_instr_detector_filter_ipppssoo_drz.fits'* and combine all exposures in a given visit by filter. Similar to standard HST products, the SVMs are drizzled to the native detector scale but are instead oriented with North up.

The second type of advanced products are referred to as Multi Visit Mosaics '*hst_skycell-pPPPPxXXyYY_instr_detector_filter_all_drz.fits'* and combine HST observations for all dates in a given detector and filter which fall within the same 0.2° x 0.2° sky cell. MVMs for ACS/WFC, WFC3/UVIS, and WFC3/IR are drizzled to a common WCS, defined by the tangent plane projection at a given reference position on the sky. These three detectors are drizzled to a common pixel scale of 0.04"/pixel with North oriented up. For WFC3/IR, a second set of drizzled images are produced at a 3x the scale at 0.12"/pixel, similar to the native resolution of the detector. These are less noisy and may be preferable to use for estimating photometry, compared to the super-sampled 0.04"/pixel layer. MVMs are not produced for the ACS/HRC or ACS/SBC detectors, since their footprints on the sky are very small, the native detector resolution is smaller than the 0.04"/pixel output grid, and relatively few datasets have successful catalog alignment.

Standard HST products and SVMs are recommended for the most accurate photometry, where the time-dependent photometry keyword values populated the image headers should be used for absolute calibration. Since the SVM products include an improved relative alignment across filters, the drizzled images are used to produce both point source and segment catalogs using small apertures. These are intended to replicate the Hubble Source Catalog, a product of the Hubble Legacy Archive, which was last updated in September 2019.

MVMs are intended to serve as 'discovery images' and may have photometric errors up to 5% for images observed over a large range of dates, due to temporal changes in the inflight sensitivity for some detectors which is not accounted for when combining each layer. For some layers, the alignment may not be optimal when different visits in a given filter are combined together, regardless of WCS type (e.g. both *a priori* and *a posteriori* solutions), and this may result in drizzled data products which appear to have multiple sets of sources. Software tools are available for users to who wish to reprocess their observations in order to improve on the current set of data products.

Acknowledgements

The development and production of Hubble Advanced products is based on the contributions of numerous talented STScI staff from the Data Management Division, including:

- 'Data Processing & Archive Services Branch' - Matthew Burger, Hien Tran, Tracy Ellis, Sara Anderson
- 'Science Instrument Calibration Software Branch'- Warren Hack, Michele de la Pena
- 'Archive Science Applications Branch' - Brian Erickson, Scott Fleming, Syed Gilani, Ben Falk, Chinwe Edeani, Clara Brasseur, Carlita Phillip
- 'Science Platform Branch'- Michael Dulude
- 'Data Analysis Tools Branch'- Mihai Cara
- 'Catalog Science Branch'- Rick White, Brian McLean, David Rodriguez, Steve Lubow
- 'Data Systems Branch'- Robert Swaters, Mike Swam, Lisa Sherbert, Hunter Brown, Lisa Gardner, Christopher Rahami, Talya Kelley, David Wolfe, Fred Romelfanger

The alignment software was developed with consultation from the STScI Astrometry Working Group members: Mike Fall (lead), Jay Anderson, Andrea Bellini, Stefano Casertano, Warren Hack, Steve Lubow, Brian McClean, David Rodriguez, Rick White, David Wolfe.

J. Mack is grateful to N. Cullings (grade 6) for helping create the RGB color mosaics shown in the Appendices. These visually striking examples highlight the power of the new MVM data products when comparing different HST detectors over a broad range of wavelengths.

The authors thank Mariarosa Marinelli, Sylvia Baggett, and Joel Green for their careful review of this report and for providing helpful suggestions to improve its content. We are also grateful to Vera Kozhurina-Platais for her detailed comments and for many useful discussions about astrometry and the correction of optical distortion in HST images. The improved absolute astrometry in these MAST products would not be possible without an accurate distortion calibration for each detector.

References

Avila. R. J. et al. 2019, *SBC Absolute Flux Calibration*, ACS ISR 2019-05

Bajaj, V. 2017, *Aligning HST images to Gaia: A Faster Mosaicking Workflow*, WFC3 ISR 2017-19

Bertin, E. & Arnouts, S. 1996, *SExtractor: Software for source extraction*", 1996A&AS..117..393B

Bradley, L. et al. 2022, astropy/photutils: 1.5.0, Zenodo. https://doi.org/10.5281/zenodo.6825092

Hoffmann, S. L. & Avila, R. J. 2017, *Updated MDRIZTAB Parameters for ACS/WFC*, ACS ISR 2017-02

Gaia Collaboration, Brown, A.G.A., et al. 2021, *Gaia Early Data Release 3. Summary of the contents and survey properties*, 2021A&A...649A..1G

Hoffmann, S. L., Mack, J., et al., 2021, *The DrizzlePac Handbook*, Version 2.0

Kozhurina-Platais, V., Borncamp, D., Anderson, J., Grogin, N., & Hack W. 2015, *ACS/WFC Revised Geometric Distortion for DrizzlePac*, ACS ISR 2015-06

Lasker, B. et al. 1990, *The Guide Star Catalog. I - Astronomical foundations and image processing*, 1990AJ...99.2019L

Lasker, B. et al. 2021, "*VizieR Online Data Catalog: The Guide Star Catalog, Version 2.4.2 (GSC2.4.2) (STScI, 2020)*", 2021yCat.1353

Lucas, R. A. & Ryon, J. E., et al., 2022, *ACS Data Handbook for Cycle 30*, Version 11.0

Mack J. & Bajaj, V., 2019, *Improved Drizzled Data Products for WFC3/IR*, WFC3 ISR 2019-05

Sahu, K. C., et al., 2021, *WFC3 Data Handbook*, Version 5.0

Skrutskie, M. F. et al. 2006, *The Two Micron All Sky Survey (2MASS)*, 2006AJ..131..1163S

Stetson, P. B. 1987, *DAOPHOT: A Computer Program for Crowded-Filed Stellar Photometry*, 1987PASP..99..191S

Whitmore et al., 2016, *Version 1 of the Hubble Source Catalog*, 2016AJ..151..134W

Appendix A: MAST Data Products for hst-skycell-p1857x08y11

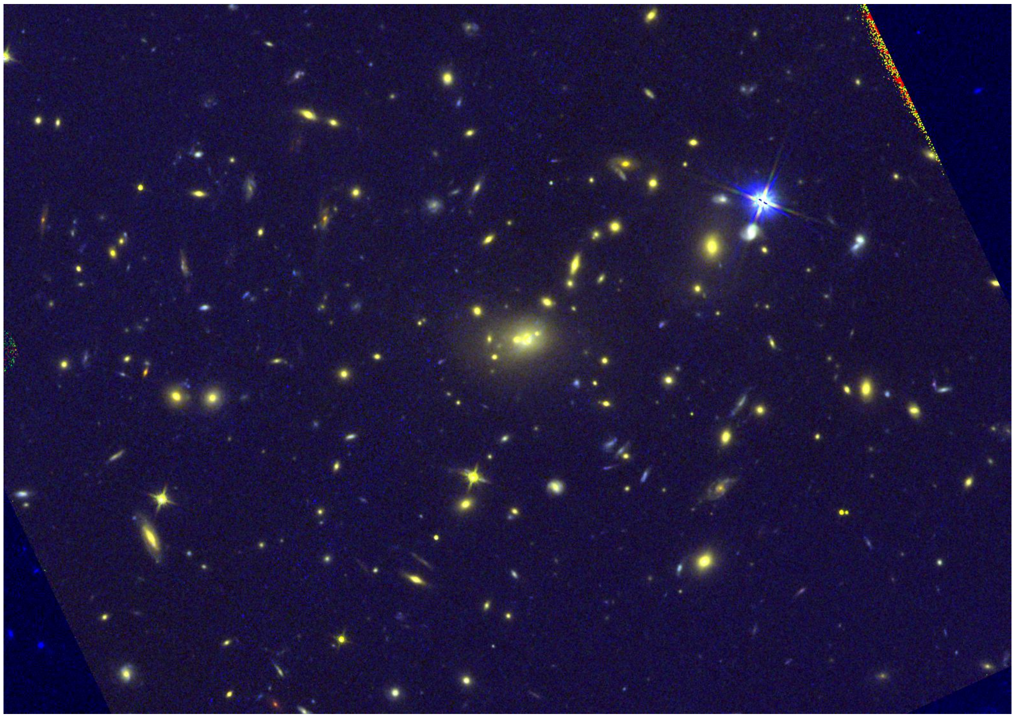

Figure A1: MVM color composite of the galaxy cluster EMACSJ0241.0+2557 in skycell-p1857x08y11 region (3.0'x2.5'). The three filter layers were derived from ACS/WFC F606W (blue), WFC3/IR F110W (green), and WFC3/IR F140W (red) images aligned to GAIAeDR3 and drizzled to the same footprint at a scale of 0.04"/pixel. (The linear feature in the upper right is an artifact due to unbonded pixels at the top edge of the WFC3/IR detector (Bushouse 2008.))

Table A1: MVM data products for skycell-p1857x08y11 shown in Figure A1 with RGB colors corresponding to the three layers. This sky cell contains the galaxy cluster EMACSJ0241.0+2557 observed in program 15312 with WFC3/IR (F110W and F140W) and in program 15843 with ACS/WFC (F606W). IR drizzled products are produced for two output scale values: a 'coarse' grid of 0.12"/pix, similar to the native plate scale, and a fine grid of 0.04"/pixel (super-sampled x3) to match the ACS/WFC drizzled data.

Filename (hst_skycell_p*-all_drz.fits)	**Prop**	**Detector**	**Filter**	**N_Exp**	**Scale ("/pix)**
hst_skycell-p1857x08y11_wfc3_ir_f140w_all_drz.fits	15132	WFC3/IR	F140W	2	0.04
hst_skycell-p1857x08y11_wfc3_ir_f140w_coarse-all_drz.fits					0.12
hst_skycell-p1857x08y11_wfc3_ir_f110w_all_drz.fits	15132	WFC3/IR	F110W	2	0.04
hst_skycell-p1857x08y11_wfc3_ir_f110w_coarse-all_drz.fits					0.12
hst_skycell-p1857x08y11_acs_wfc_f606w_all_drc.fits	15843	ACS/WFC	F606W	3	0.04

Table A2: SVM data products for skycell-p1857x08y11 shown in Figure A1. These include calibrated *'_fl[tc].fits'* and drizzled *'_dr[zc].fits'* images with improved relative alignment across filters in a given visit compared to standard products. Point and segment photometry catalogs are computed for each detector+filter based on a combined detection image *'_total_dr[zc].fits'*. The exposure-level calibrated images *'_flt.fits'* (WFC3/IR) and the CTE-corrected *'_flc.fits'* images (ACS/WFC) are used to produce the MVM products in Table A1, with no additional alignment. The MAST 'minimum recommended' data products are shown in bold.

Filename hst_propid_obsetid_instr_detector_filter_ipppss_*.*	**TYPE**	**Rootname**
hst_15132_0l_wfc3_ir_f140w_idkl0l_drz.fits	DRZ F140W	idkl0l
hst_15132_0l_wfc3_ir_f140w_idkl0l_point-cat.ecsv	Point Source Catalog	idkl0l
hst_15132_0l_wfc3_ir_f140w_idkl0l_segment-cat.ecsv	Segment Catalog	idkl0l
hst_15132_0l_wfc3_ir_f140w_idkl0ldo_flt.fits	F140W FLT Exp 1	idkl0ldo
hst_15132_0l_wfc3_ir_f140w_idkl0ldo_drz.fits	F140W DRZ Exp 1	idkl0ldo
hst_15132_0l_wfc3_ir_f140w_idkl0ldm_flt.fits	F140W FLT Exp 2	idkl0ldm
hst_15132_0l_wfc3_ir_f140w_idkl0ldm_drz.fits	F140W DRZ Exp 2	idkl0ldm
hst_15132_0l_wfc3_ir_f110w_idkl0l_drz.fits	DRZ F110W	idkl0l
hst_15132_0l_wfc3_ir_f110w_idkl0l_point-cat.ecsv	Point Source Catalog	idkl0l
hst_15132_0l_wfc3_ir_f110w_idkl0l_segment-cat.ecsv	Segment Catalog	idkl0l
hst_15132_0l_wfc3_ir_f110w_idkl0ldk_flt.fits	F110W FLT Exp 1	idkl0ldk
hst_15132_0l_wfc3_ir_f110w_idkl0ldk_drz.fits	F110W DRZ Exp 1	idkl0ldk
hst_15132_0l_wfc3_ir_f110w_idkl0ldj_flt.fits	F110W FLT Exp 2	idkl0ldj
hst_15132_0l_wfc3_ir_f110w_idkl0ldj_drz.fits	F110W DRZ Exp 2	idkl0ldj
hst_15843_60_acs_wfc_f606w_je2860_drc.fits	DRZ F606W	je2860
hst_15843_60_acs_wfc_f606w_je2860_point-cat.ecsv	Point Source Catalog	je2860
hst_15843_60_acs_wfc_f606w_je2860_segment-cat.ecsv	Segment Catalog	je2860
hst_15843_60_acs_wfc_f606w_je2860ju_flc.fits	F606W FLC Exp 1	je2860ju
hst_15843_60_acs_wfc_f606w_je2860ju_drc.fits	F606W DRC Exp 1	je2860ju
hst_15843_60_acs_wfc_f606w_je2860jv_flc.fits	F606W FLC Exp 2	je2860jv
hst_15843_60_acs_wfc_f606w_je2860jv_drc.fits	F606W DRC Exp 2	je2860jv
hst_15843_60_acs_wfc_f606w_je2860jx_flc.fits	F606W DRC Exp 3	je2860jx
hst_15843_60_acs_wfc_f606w_je2860jx_drc.fits	F606W FLC Exp 3	je2860jx
hst_15132_0l_wfc3_ir_total_idkl0l_drz.fits	Detection Image: IR	idkl0l
hst_15132_0l_wfc3_ir_total_idkl0l_point-cat.ecsv	Point Source Catalog	idkl0l
hst_15132_0l_wfc3_ir_total_idkl0l_segment-cat.ecsv	Segment Catalog	idkl0l
hst_15843_60_acs_wfc_total_je2860_drc.fits	Detection Image: ACS	je2860
hst_15843_60_acs_wfc_total_je2860_point-cat.ecsv	Point Source Catalog	je2860
hst_15843_60_acs_wfc_total_je2860_segment-cat.ecsv	Segment Catalog	je2860

Table A3: Standard data products for skycell-p1857x08y11 shown in Figure A1. These include the drizzle-combined images *'_dr[cz].fits'* for each filter (with and without CTE-correction), the exposure-level calibrated images *'_fl[tc].fits'*, the WFC3/IR intermediate products *'_ima.fits'* with no ramp fitting, and the raw *'_raw.fits'* images. The calibrated images *'_fl[tc].fits'* are used as input for the SVM data products in Table A2. The MAST 'minimum recommended' data products are shown in bold.

Proposal	**Filename** (ipppssoot_*.fits)	**TYPE**	**ObsetID** (Visit)	**Rootname**
15132	**idkl0l020_drz.fits**	F140W DRZ	01	idkl0l020
	idkl0ldmq_flt.fits	F140W FLT Exp 1		idkl0ldmq
	idkl0ldkm_ima.fits	F140W IMA Exp 1		idkl0ldmq
	idkl0ldkm_raw.fits	F140W RAW Exp 1		idkl0ldmq
	idkl0ldoq_flt.fits	F140W FLT Exp 2		idkl0ldoq
	idkl0ldoq_ima.fits	F140W IMA Exp 2		idkl0ldoq
	idkl0ldoq_raw.fits	F140W RAW Exp 2		idkl0ldoq
15132	**idkl0l010_drz.fits**	F110W DRZ	01	idkl0l010
	idkl0ldjq_flt.fits	F110W FLT Exp 1		idkl0ldjq
	idkl0ldjq_ima.fits	F110W IMA Exp 1		idkl0ldjq
	idkl0ldjq_raw.fits	F110W RAW Exp 1		idkl0ldjq
	idkl0ldkq_flt.fits	F110W FLT Exp 2		idkl0ldkq
	idkl0ldkq_ima.fits	F110W IMA Exp 2		idkl0ldkq
	idkl0ldkq_raw.fits	F110W RAW Exp 2		idkl0ldkq
15843	**je2860010_drc.fits**	F606W DRC	60	je2860010
	je2860010_drz.fits	F606W DRZ		je2860010
	je2860juq_flc.fits	F606W FLC Exp 1		je2860juq
	je2860juq_flt.fits	F606W FLT Exp 1		je2860juq
	je2860juq_raw.fits	F606W RAW Exp 1		je2860juq
	je2860jvq_flc.fits	F606W FLC Exp 2		je2860jvq
	je2860jvq_flt.fits	F606W FLT Exp 2		je2860jvq
	je2860jvq_raw.fits	F606W RAW Exp 2		je2860jvq
	je2860jxq_flc.fits	F606W FLC Exp 3		je2860jxq
	je2860jxq_flt.fits	F606W FLT Exp 3		je2860jxq
	je2860jxq_raw.fits	F606W RAW Exp 3		je2860jxq

Appendix B: Sample MVM Color Composites

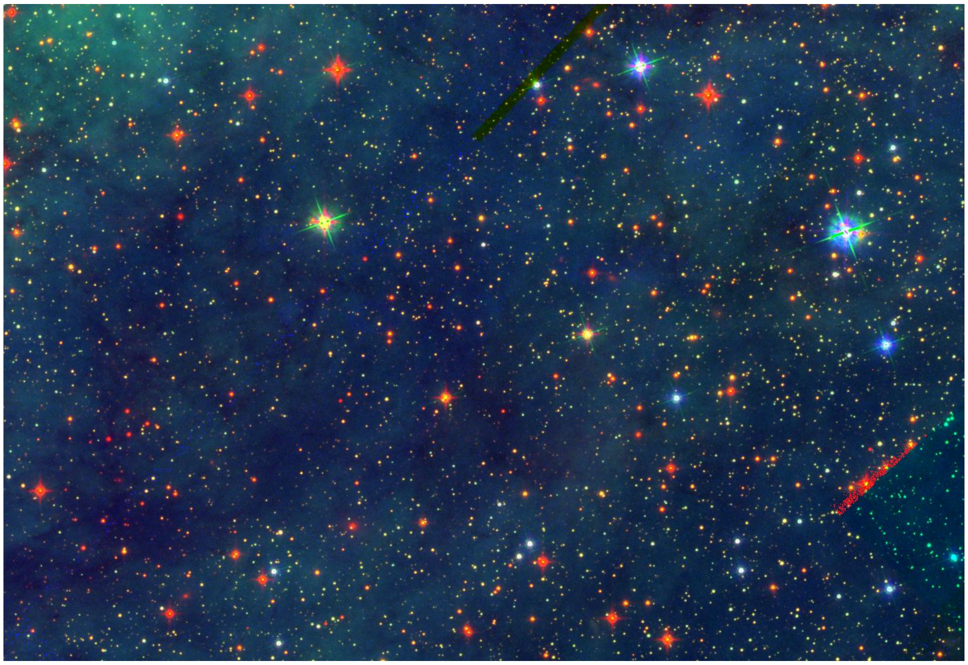

Figure B1: MVM color composite of hst_skycell-p0081x16y15 for a small region (2'x1.3') in 30 Doradus which details the quality of the alignment for three different detectors. The WFC3/UVIS F336W layer is shown in blue, the ACS/WFC F555W layer in green, and the WFC3/IR F110W filter in red. (The red feature at lower-right is an artifact in WFC3/IR, see Figure A1.)

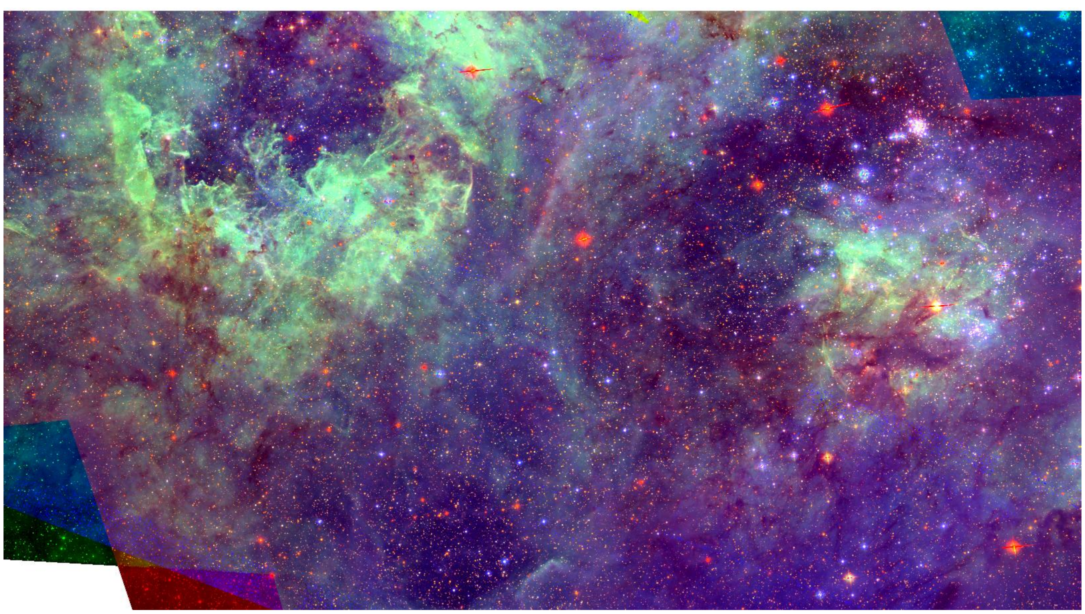

Figure B2: MVM color composite of hst_skycell-p0081x15y14 for much larger region (8.5'x4.5') of 30 Doradus. The WFC3/UVIS F336W layer is shown in blue, the ACS/WFC F555W layer in green, and the ACS/WFC F775W layer in red.

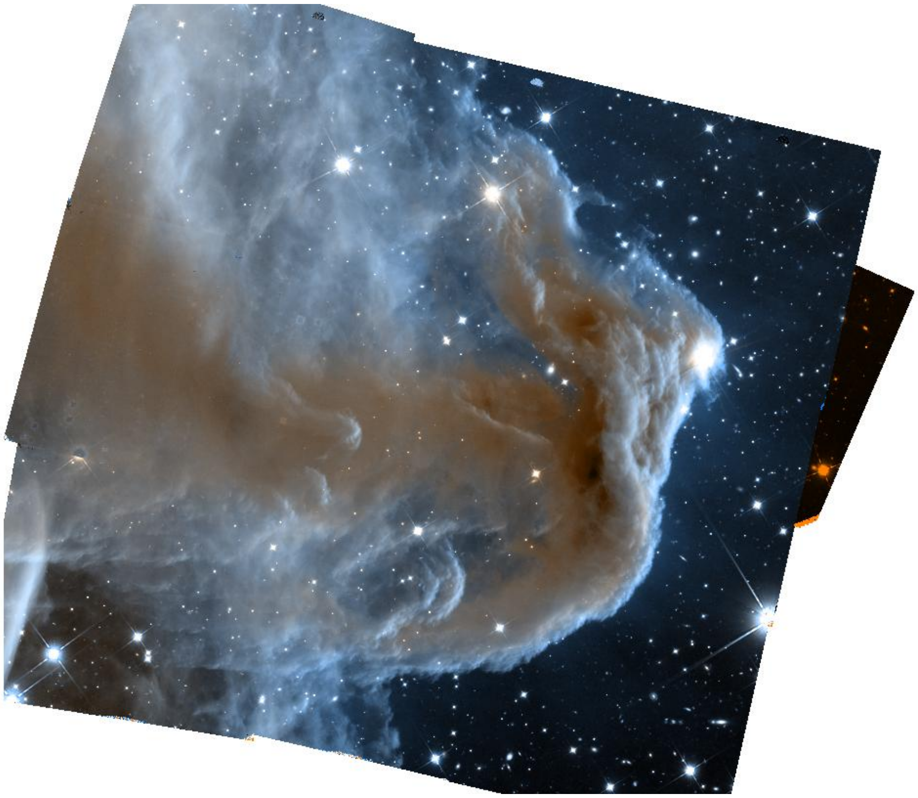

Figure B3: MVM color composite of hst_skycell-p1253x06y09 showing a 3x3 mosaic of the Horsehead Nebula in WFC3/IR. The F110W layer is shown in blue and the F160W layer in red.

Figure B4: MVM color composite of hst_skycell-p1101x11y04 showing a 2x2 mosaic of the Sombrero Galaxy in ACS/WFC with F435W (blue layer), F555W (green), and F625W (red).

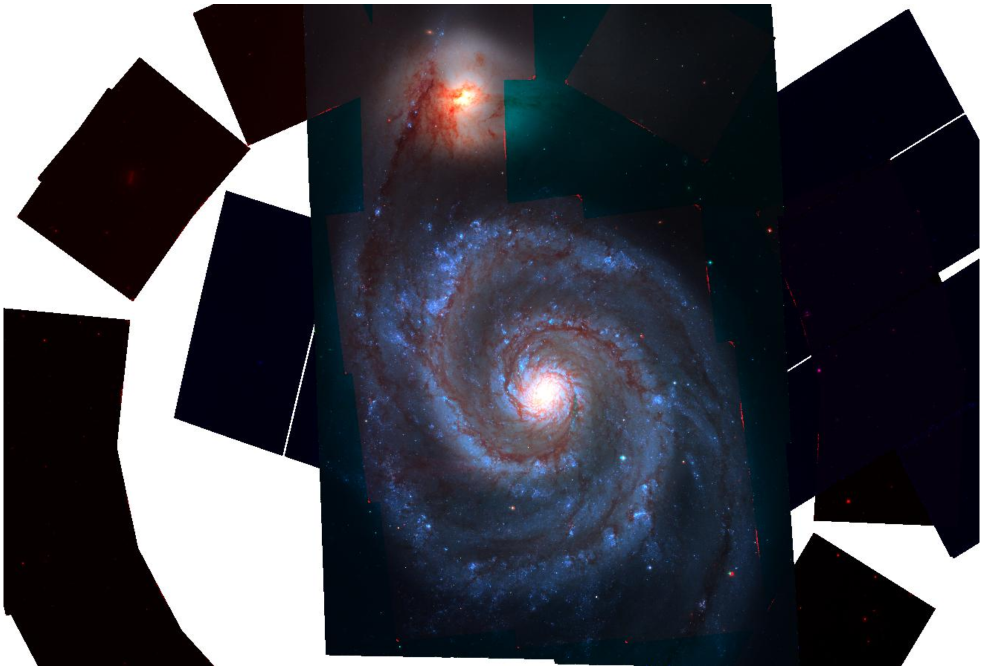

Figure B5: MVM color composite of skycell-p2261x11y16 for a large region surrounding M51. The ACS/WFC F435W layer is shown in blue, the ACS/WFC F555W layer in green, and the WFC3/IR F110W layer in red. A large number of parallel WFC3/IR observations sweep out a circle around the target and correspond to a range of dates from Oct 2016 – Sep 2017 when observing with ACS/WFC as the prime instrument in program 14704. (Not all regions of the field were observed in each of the three filters in this composite image.)

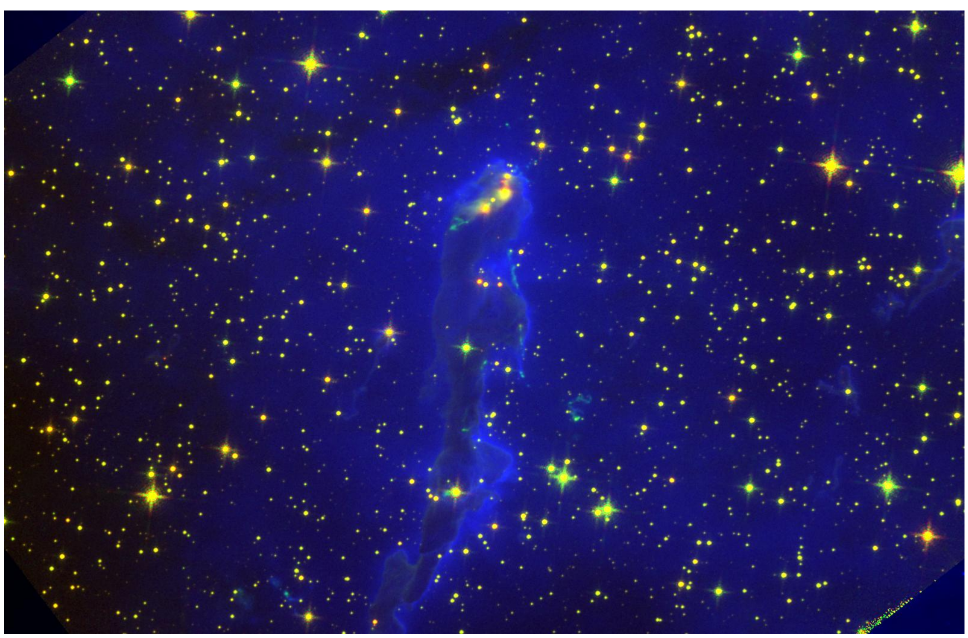

Figure B6: MVM color composite of hst_skycell-p0165x08y19 for a small region (2'x1.3') in the Carina Nebula. The ACS/WFC F658N layer is shown in blue, the WFC3/IR F126N layer in green, and the WFC3/IR F167N layer in red.