



STScI | SPACE TELESCOPE
SCIENCE INSTITUTE

Instrument Science Report ACS 2024-02

The impact of CTE on point source detection in simulated ACS/WFC imaging data

David V. Stark, Norman Grogin

May 14, 2024

ABSTRACT

*Using mock observations, we investigate how imperfect charge transfer efficiency (CTE) affects the detection of point sources in imaging data from the Advanced Camera for Surveys Wide Field Channel (ACS/WFC). Using two public source detection algorithms, **hst1pass** and **DAOStarFinder**, we analyze the recovery of sources in both FLT and (CTE-corrected) FLC images, as well as an identical set of idealized images with perfect CTE. The simulations incorporate the expected worsening of CTE losses as a function of ACS lifetime and exposure background level. Optimal signal-to-noise ratio (SNR) for sources with brightnesses less than a few hundred electrons is obtained at backgrounds of $\sim 30e^-/\text{pixel}$, where there is a balance between decreasing CTE losses and increasing background Poisson noise. Consistently, the sensitivity of FLC/FLT images is maximized at these sky levels when using **hst1pass** for detection, although this trend is less apparent when using **DAOStarFinder**. Individual FLT exposures are slightly more sensitive than FLCs at low backgrounds, but this behavior generally reverses when multiple exposures are combined. Importantly, FLC exposures do not match the depth and SNR of idealized exposures with perfect CTE until backgrounds reach $> 100e^-/\text{pixel}$. False detections become an issue when relaxing the allowed point-spread function goodness-of-fit in **hst1pass** although the false sources do not appear to be linked to CTE or its correction. False detections are very common in **DAOStarFinder** output and do appear to be associated with either the detection of CTE trails or amplified noise from the CTE-correction. However, simple cuts on source properties (SNR, sharpness) can reject most of these false detections.*

1 Introduction

As the Advanced Camera for Surveys (ACS) aboard the Hubble Space Telescope (HST) ages in low-Earth orbit, it is continuously bombarded by high energy particles that damage its CCD detectors. This damage creates charge traps that reduce the Charge Transfer Efficiency (CTE) of the CCDs, i.e., the ability of collected electrons in a given pixel to successfully transfer in the parallel and serial directions during readout. Degraded CTE leads to “trails” behind bright sources caused by the delayed readout of some of their electrons, leading to lower signal-to-noise (SNR) and possibly shifts in the derived positions of some sources. Generally, degraded CTE has an increasing impact on data quality as source brightness and backgrounds decrease, and as the number of transfers for readout increases.

The CTE of ACS has been extensively studied and tools have been developed to correct for it. The ACS team has provided software to apply “photometric” and “pixel-based” CTE corrections. The former corrects point source fluxes using a curvilinear function depending on source flux, background level, and date (Chiaberge, 2012; Chiaberge & Ryon, 2022). The pixel-based method uses a forward modeling technique to correct the flux of individual pixels, redistributing it back to where it “should” be (details can be found in Anderson & Bedin, 2010; Anderson & Ryon, 2018). This latter approach has been incorporated into the standard ACS pipeline and is applied to all exposures to generate FLC images. Efforts are also underway to quantify and correct for the smaller, but detectable, impact of imperfect serial-CTE (xCTE).

For sufficiently bright sources, degraded CTE will not affect one’s ability to detect those sources in an image. More generally, if one knows source positions a-priori, observations can be tuned to place those sources near the CCD serial registers, and/or one can apply the photometric CTE correction to obtain an accurate estimate of the original flux. However, the situation is different if the positions of sources, especially fainter ones, are unknown ahead of time. The degradation of SNR will make it more difficult to blindly detect sources far from the CCD serial registers, so source recovery will not be uniform over a given exposure. While pixel-based CTE corrections mitigate this issue to some extent, they are not effective on sources with amplitudes close to the read noise. Historically, HST has proven to be an invaluable tool to discover new objects (the Hubble Deep Field is an excellent example of this use case; Williams et al. 1996). Therefore, it is crucial to understand exactly how CTE affects our ability to detect sources blindly. Once we fully understand the limitations imposed by imperfect CTE, observations and/or analyses can be adjusted to account for them.

One benefit of developing the pixel-based CTE correction algorithm is the ability to model the impact of CTE losses. Using this tool, we can generate realistic simulated images that yield large-number statistics and highly controlled parameters, allowing a clear view of the impact of imperfect CTE. In this report, we utilize such mock data sets to understand how CTE drives the relationship between image background and point source signal-to-noise ratio (SNR), as well as point-source detectability. We also analyze if CTE can introduce false sources during source detection, and how these can be avoided.

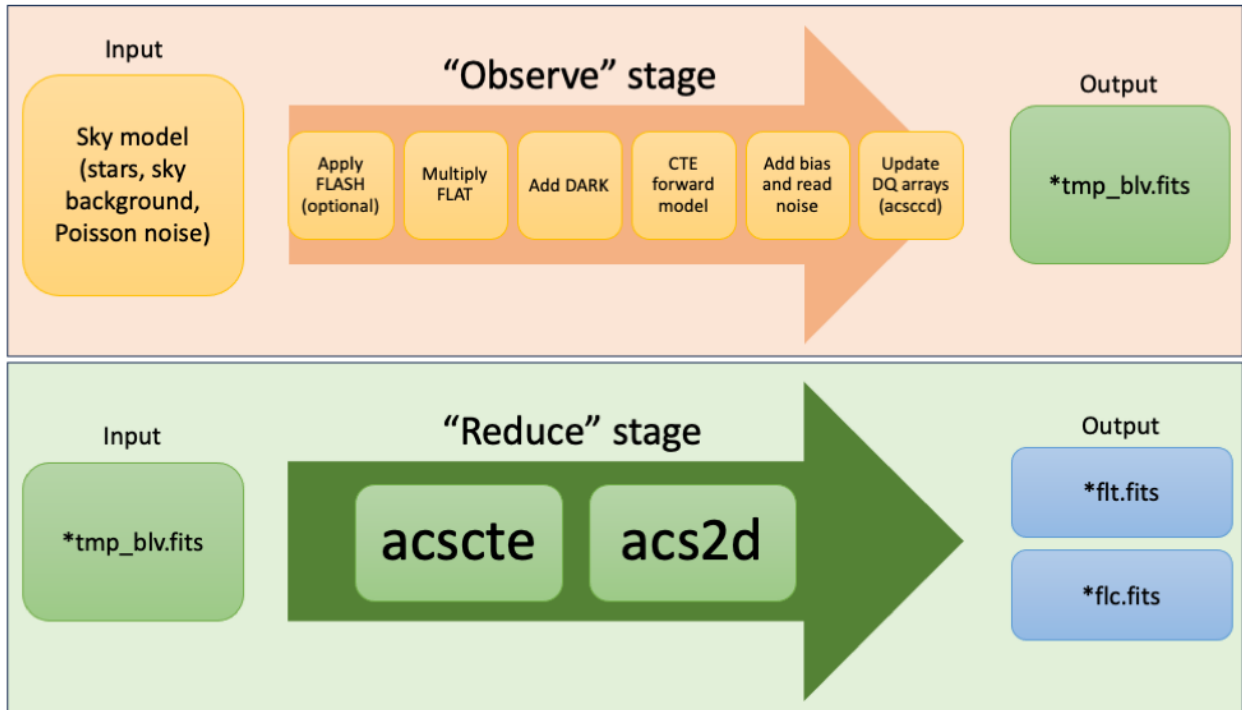


Figure 1: A summary of the pipeline for making mock ACS/WFC observations and reductions.

2 Methods

In this section, we describe the pipeline for generating mock images (Section 2.1), the properties of simulated images used in this study (Section 2.2), and the methods for source detection (Section 2.3).

2.1 Simulating ACS/WFC Imaging Data

To generate mock ACS/WFC images, we utilize an algorithm that involves (a) creating a sky model, (b) applying all observational effects to it, and (c) passing it through the standard ACS reduction pipeline. A visual summary of the approach is shown in Figure 1. The methodology used here is largely based on a Jupyter-notebook¹ provided by the ACS team describing how to implement the CTE forward modeler, but reworked into a stand-alone program that allows us to more automatically generate large numbers of simulated images. In the following sections, we outline the key steps in this process.

¹https://github.com/spacetelescope/acs-notebook/blob/master/acs_cte_forward_model/acs_cte_forward_model_example.ipynb

2.1.1 Initializing the sky model

We start by selecting a `RAW` file from an existing observation with an observation date matching as closely as possible to the desired date of our mock exposure. The date plays an important role of setting the lifetime of ACS for the simulation, which will in turn affect the relative impact of imperfect CTE. We then modify any additional header keywords as needed, most importantly the exposure time and filter. While post-flash is not typically performed on ACS/WFC imaging data², if we desire to simulate its impact on our mock data, header keywords for the flash reference file (`FLSHFILE`) and flash duration (`FLASHDUR`) are updated, while the flash current (`FLASHCUR`) keyword is set to `ON` and the flash status keyword (`FLASHSTA`) is set to `SUCCESSFUL`. Other aspects of the raw image (especially its data HDUs) will have no impact on the observation or will be overwritten in subsequent steps. While in principle we could create a new `RAW` image file from scratch, using an existing file and overwriting the relevant details is a useful shortcut. In particular, the various reference files (for dark correction, flat fielding, etc.) are based on the date and already populated. Once the parameters of the `RAW` have been modified, it is passed through `acsccd` to create a `blv_tmp.fits` file that has the overscan regions removed and bias subtracted. The science data arrays are then completely erased.

With the now empty `blv_tmp` file, we first add the desired uniform sky level and its corresponding Poisson noise. Next, stars of preset brightness and position are added to the image. Our purpose is not to create a model with a realistic spatial distribution of stars, but rather one that allows us to reliably characterize the impact of CTE losses. Therefore, we arrange the stars on each WFC chip in a regular grid with the stars separated by 64 pixels in the x direction and 32 pixels in the y direction. To ensure any CTE trails do not overlap stars in neighboring rows, the x positions of stars in each row are shifted by 16 pixels relative to the previous row. There are a total of 2032 stars per chip, or 4064 stars per image.

To create the actual point sources, we use position-dependent PSF models described in Anderson (2016). We interpolate these PSFs as a function of position for every model star, but note that all stars are always centered on individual pixels. In other words, this report does not investigate how CTE losses are related to different pixel phase sampling, which will impact the profiles of individual stars on the CCD depending on where they are centered relative to pixel centers. Star brightnesses are randomly assigned using logarithmic sampling between the range of brightnesses we desire (described below), and Poisson noise is added to the stars. An example sky model is shown in Figure 2.

In principle, one can add additional features such as extended sources, or additional artifacts like cosmic rays or satellite trails. For the purposes of this work, we do not include these additional features.

2.1.2 Observational Effects

With the sky model in place, we apply the various observational artifacts that all real ACS/WFC images are subjected to:

²Users are typically advised to increase exposure times to raise backgrounds, or place their target near the serial registers, before resorting to using post-flash (Ogaz et al., 2014; Miles, 2018).

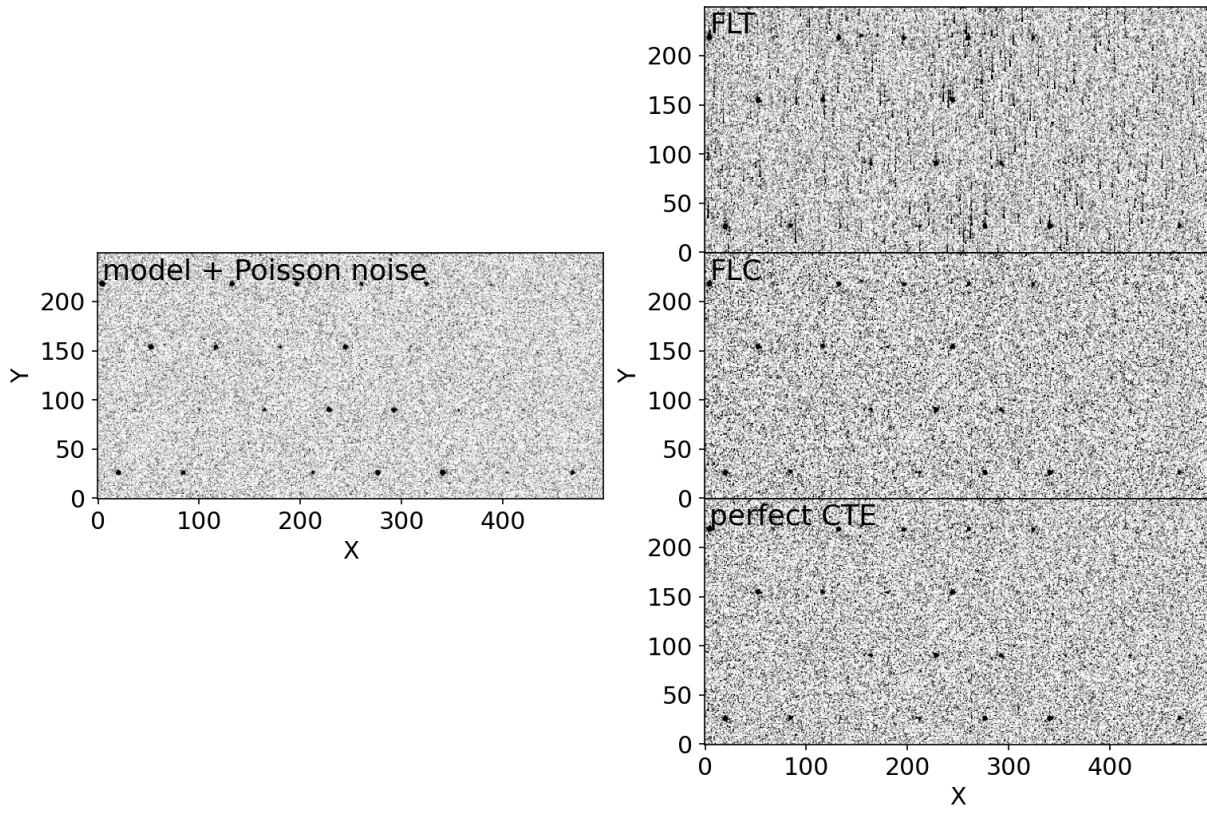


Figure 2: Zoomed in cutouts of an example input sky model fed into the simulation pipeline (left) and the resulting output FLT (top right), FLC (middle right), and perfect-CTE (bottom right) images. This specific example is taken from simulation set 1 and shows a mock 2023 350s second exposure.

- Post-flash: If desired, post-Flash is applied to the data by obtaining the reference flash file (`FLSHFILE`) and scaling it by the specified flash duration. This scaled image, with Poisson noise added, is then applied to the model image.
- Pixel-to-pixel sensitivity: To simulate the non-uniform sensitivity of pixels in the CCDs, we multiply the model image by the reference flat field (`PFLTFILE`).
- Dark Current: To mimic the accumulation of dark current during the exposure, we multiply the reference dark image (`DRKCFILE`) by the appropriate time (this includes the nominal exposure time, the flash duration, plus any overheads), apply Poisson noise to the scaled dark image, then add it to the model image.
- Charge Transfer Efficiency: One of the most important parts of our pipeline is simulating the impact of CTE. We use the same CTE forward modeler used to correct real ACS/WFC images to create the CTE effects. Details of the model are presented in Anderson & Bedin (2010) and Anderson & Ryon (2018). This step only models the impact of imperfect CTE in the parallel (y) direction, not in the serial (x) direction.
- Read noise: We add amplifier-dependent read noise to each image. The noise follows a Gaussian distribution with a mean of zero and a standard deviation taken from the corresponding header keywords (`READNSE[A-D]`).

2.1.3 Updating the ERR and DQ extensions

We now have a model image that is equivalent to a realistic `blv_tmp` image, but the `ERR` and `DQ` arrays have not been modified from the original input. The error arrays are recalculated by adding in quadrature the noise terms from the (a) sky model including various observational effects, (b) the read noise (taken from the image header), and (c) the bias error (determined from the superbias reference file, `BIASFILE`).

The `DQ` arrays are repopulated by rerunning `acsccd`, after first renaming the `blv_tmp` file to have the suffix `raw.fits` (which just allows `acsccd` to run properly). Prior to running `acsccd`, we set the `DQICORR` and `SINKCORR` header keywords to `PERFORM`. This stage regenerates the `DQ` extension from scratch by first flagging known permanently bad pixels. Subsequently, sink pixels, which are determined from a reference file generated after each anneal, are flagged. Re-running `acsccd` results in a file that again has the suffix `blv_tmp`. We then load the superbias `DQ` extensions and combine them with the model image `DQ` extensions. Note that we do not explicitly simulate permanent CCD defects like sink pixels, but by flagging them in the `DQ` array, we can ensure consistency in flagging data in simulations and real observations. A side effect of rerunning `acsccd` is that it assumes the images are in units of DN, so it multiplies the data arrays by the gain. However, our model is already in units of electrons, so we simply divide out the gain for each amplifier (the gain values are taken from the header).

2.1.4 Re-running the reduction pipeline

At this stage, we have a model `blv_tmp` file equivalent to what would be generated from a real observation of our model sky after processing its `RAW` file with `acsccd`. We finish the

simulation by passing this model image through the standard reduction tasks, first running `acscte` which applies the pixel-based CTE correction, and then running `acs2d` which applies the dark, flatfield, and (optionally) post-flash corrections. These steps result in the FLT and (CTE-corrected) FLC files used for our analysis.

2.2 The full suite of simulations

Using the pipeline described in the previous section, we generate several sets of simulations for different purposes:

- Set 1: This simulation set contains background levels of $10e^-/\text{pix}$, $20e^-/\text{pix}$, and $30e^-/\text{pix}$ with corresponding exposure times of 120s, 240s, and 350s respectively. Point source brightnesses range from $\sim 30e^-$ to $\sim 3000e^-$. Identical simulations are generated for the years 2003, 2013, and 2023. This simulation set is primarily used to study the time-dependent impact of CTE losses. ACS was installed in 2002, so 2003 reflects near-pristine CTE, while 2023 reflects the latest (at the time of writing this report) ACS data, and 2013 serves as an intermediate sampling.
- Set 2: These simulations mimic 2023 data and set a very low background level of $5e^-/\text{pix}$ (60s exposure). Backgrounds are elevated by adding progressively longer post-flashing up to 14.3 seconds, yielding final local backgrounds (measured around each source) ranging from 5 to $200 e^-/\text{pix}$. Source fluxes range from $\sim 10e^-$ to $150e^-$. These simulations are used to test SNR as a function of sky background for very faint sources in the presence of CTE losses.
- Set 3: This simulation set is the same as Set 2, except source brightnesses range from $\sim 30e^-$ to $\sim 3000e^-$. These simulations are used to test SNR and survey completeness limits in the presence of CTE losses.

All simulations generate stars assuming use of the F775W filter³, although we do not expect the choice of filter to significantly change the results of our analysis. Each unique simulation configuration has ten different versions, each with its own set of random star brightnesses. In other words, for a given configuration (e.g., 2013 mock, background level of $20e^-/\text{pix}$, no flash) there are 4064 stars per image, and with ten unique versions, there are 40,640 unique stars per simulation configuration. The same ten iterations of random star brightnesses are reused for each different simulation configuration. This consistency allows for a fair comparison in the performance of different simulated configurations.

We also generate identical versions of each simulation with perfect CTE, which is accomplished by bypassing the CTE modeling and CTE correction stages in the simulation pipeline. Creating these idealized images enables us to clearly understand how CTE is changing point source SNR and detectability beyond the expected behavior from typical noise sources. Examples of mock FLT, FLC, and perfect-CTE images are shown in Figure 2.

³This filter was used by random chance based on the filter used in one of the first simulations generated.

2.3 Source Detection

There are numerous publicly available codes for source detection. It is beyond the goals of this report to thoroughly review the performance of each on CTE-impacted data sets. Instead, we look at two source detection codes that operate in distinct ways, and the results of our analysis should generally be relevant for users employing other detection algorithms.

Source detection is primarily carried out using `hst1pass` that performs PSF-fitting photometry on undersampled HST images like those from ACS/WFC (Anderson, 2022). We employ the code with a minimum isolation (`HMIN`) of 5 pixels, a minimum central flux (`FMIN`) equal to the median sky value plus five times the its standard deviation (using outlier rejected median and standard deviations) measured across the full pair of WFC CCDs, and position-dependent ePSFs. All other parameters are left at their default values.

One of the measurements `hst1pass` returns is a PSF “goodness-of-fit” parameter, q , for each source with a peak flux above `FMIN`. By default, `hst1pass` enforces a maximum q of 0.5, rejecting any sources that do not satisfy this condition. We adopt this default, but will discuss in Section 4.1 how allowing higher values of q can impact source recovery and purity. We also note that `hst1pass` does not set a lower limit on q , or check whether a resulting detected source has negative total flux. In the course of this project we discovered numerous false “detections” with $q \ll 0$, which were deemed to be regions with a single pixel above the detection threshold but PSF fits that yielded negative integrated flux. We found that $> 99.9\%$ of $q < 0$ sources were false, while $> 99.9\%$ of real sources had $q > 0$. Therefore, we simply filtered out anything with $q < 0$ for the remainder of this project.⁴

The second source detection code we use for this study is `DAOStarFinder` included with the `photutils` Python package⁵. This routine works by smoothing an image by a Gaussian kernel, then finding all pixels with values above a certain threshold (which we again set to the median sky value plus five times the sigma-clipped standard deviation of the sky). We use a Gaussian kernel with full width half maximum (FWHM) of 2.5 pixels, roughly equal to the WFC F775W PSF FWHM. We also set a minimum separation between sources of five pixels to match the isolation criterion from `hst1pass`. `DAOStarFinder` does not employ PSF fitting, but does apply some basic shape constraints which we leave at the defaults ($0.2 < \text{sharpness} < 1.0$, $-1 < \text{roundness} < 1$). Not modifying the allowed range of shape properties means the catalogs contain false detections, which will be discussed further in Section 4.2.

3 Results

In the following section, we present our results on how source SNR, source detectability, and the presence of false detections is related to imperfect CTE. Note that for our analyses of source detectability, we do not modify the source catalogs returned by `hst1pass` or `DAOStarFinder` in any way. In reality, there may be alternative configurations for these programs that can increase source recovery, or selections on the catalogs that can elimi-

⁴The $q < 0$ sources represent unintended behavior and will be removed in future versions of `hst1pass` (J. Anderson, private communication).

⁵<https://photutils.readthedocs.io/en/stable/api/photutils.detection.DAOStarFinder.html>

nate false detections (although these selections may also remove real sources). Any of these choices may impact the specific brightness completeness limits we present below. However, the *relative* trends between source brightness, background level, and detectability are what we emphasize.

3.1 Historical Trends in Source Recovery

The first analysis we conduct using our mock data sets is to broadly investigate source recovery as a function of source intrinsic brightness, background level, and year. Prior observational studies have established that CTE losses are worse at low background levels and faint sources, and have been progressively worsening since ACS was installed aboard HST in 2002. As a result, we expect source recovery to decline with time and increasing sky background relative to model images with perfect CTE. For this analysis, we use simulation set 1.

The source recovery rates using `hst1pass` are plotted in Figure 3, which limits the analysis to the quadrant of the detector furthest from the serial registers where CTE-losses are most severe. These figures confirm the expected trends, but additional useful insight can be gained. First, the FLC images, which are “corrected” for CTE-losses, do not recover sources down to the same level as an image with perfect CTE. In fact, the recovery rate is only marginally better than for FLT images that are uncorrected. Second, at very low backgrounds ($\sim 10e^-$), the recovery of FLC images can actually be *worse* than that of FLT images. We discuss this behavior and possible regions for it further in Section 3.3.

3.2 Background levels to optimize signal-to-noise

Considering only the impact of CTE, a target with some fixed intrinsic brightness in electrons will have higher SNR with increasing background, which occurs because trailing is decreased when a uniform global background fills a portion of CCD pixels’ charge traps. In reality, a higher sky background comes with a higher noise contribution, which serves to lower the SNR of the target. In principle, there should be an optimal background level that finds a balance between these two competing effects and maximizes the SNR.

To investigate this question, we consider a somewhat extreme case of a 60s observation that yields an extremely low sky background of $\sim 5e^-$. In such a regime, CTE losses should be severe, but we can progressively increase sky levels by post-flashing the data. The expected behavior in this situation is that the SNR will increase with background, and at some point reverse and decline as the contribution of sky noise overtakes the preservation of flux from reduced CTE losses. This behavior can be seen qualitatively in Figure 4 which estimates SNR using Equation 1 in Section 9.3.1 of the ACS Instrument Handbook (Ryon & Stark, 2023), but incorporating the empirically determined losses due to CTE from Equation 2 of Chiaberge & Ryon (2022) into the source count rate, assuming the source needs 2000 y transfers to be read out. Readers should be aware that the trends in Figure 4 extrapolate the results of Chiaberge & Ryon (2022) to source fluxes below where the method was calibrated ($< 250e^-$), so the results should be treated as qualitative only.

In an attempt to more reliably determine the optimal background level, we use simulation sets 2 and 3 and directly measure the SNR at all known source positions in the FLT, FLC, and

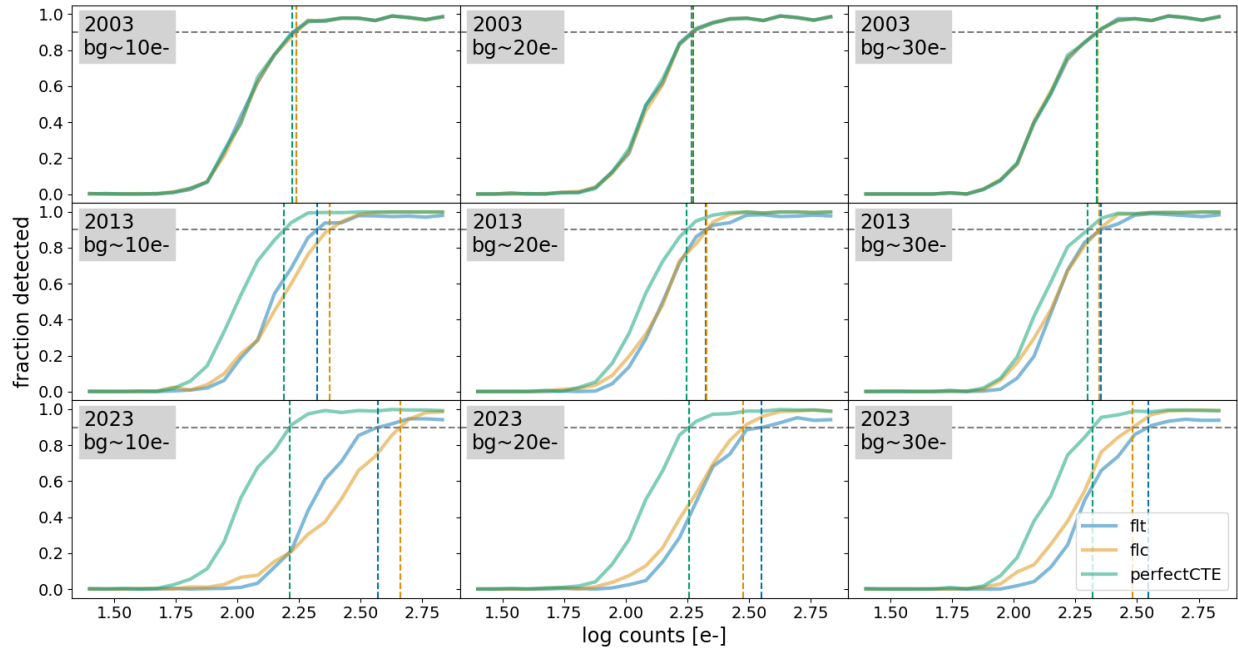


Figure 3: The detection rate of mock point sources as a function of intrinsic brightness, background level, and year. The horizontal dashed line indicates where 90% of sources are detected, and the vertical dashed lines show the brightness where each curve crosses the 90% detection threshold. Note that the different lines in the 2003 mock almost entirely overlap causing them to be hard to distinguish. The results shown here are limited to the quadrant of each detector furthest from the CCD serial registers where CTE losses are most pronounced. Our simulation pipeline recreates expected trends where losses due to CTE increase with time since ACS installation and decreasing background level, leading to lower overall sensitivity.




Figure 4: Theoretical curves of SNR versus background for stars of different intrinsic brightness, f (indicated in the legend). These curves are obtained by combining Equation 1 for SNR from Section 9.3.1 of the ACS Instrument Handbook (Ryon & Stark, 2023) with Equation 2 for losses due to CTE from Chiaberge & Ryon (2022). These curves assume use of a three pixel circular aperture to measure brightness. The vertical dashed line indicates the recommended minimum background level of $30e^-/\text{pix}$. Note that the Chiaberge & Ryon (2022) formula is being applied to stars with brightnesses below where it was calibrated ($< 250e^-$), so the results are qualitative only.

perfect-CTE images. We then characterize the mean SNR as a function of source intrinsic brightness and locally measured background. The results of this analysis are shown in Figure 5, which is limited to sources in the CCD quadrants furthest from the serial registers.

For fainter sources below a few hundred counts, we observe a rapid increase followed by a gradual decline in SNR as a function of background level. The maximum SNR occurs around a background of $\sim 30e^-/\text{pix}$, which is the current suggested minimum sky background (Ryon & Stark, 2023). For brighter sources, severe CTE losses are not evident at low backgrounds, and we only observe the gradual decline in SNR as background level increases. However, the SNRs in both the FLT and FLC images are often slightly lower than the SNRs from the perfect-CTE images.

3.3 Present-day source recovery vs background

Given the dependence of SNR on background demonstrated in Section 3.2, we expect blind detection of sources, which is typically dependent on a SNR threshold, to be impacted in a similar fashion. Figure 6 demonstrates the recovery of point sources with `hst1pass` as a function of their brightness and the locally measured background in FLT, FLC, and perfect-CTE images. For this analysis, we use simulation set 3. To highlight the impact of CTE, the data in Figure 6 are limited to the quadrant of each WFC CCD furthest from the CCD serial registers. To test whether a source is recovered, we compare the true positions and the detected star positions. If a true star position has a match within two pixels, we call it recovered. This separation of two pixels is motivated by the fact that CTE can lead a source centroid to move by spreading flux out in the y direction. We adopt the default maximum PSF goodness-of-fit of $q = 0.5$. Raising this threshold could result in a higher detection rate of stars (at the expense of false detections; Section 4.1). However, we do not expect any qualitative trends to change.

The results for perfect-CTE images, provided for reference, show the expected decreasing sensitivity with increasing background level, but sharply different behavior is observed in the FLTs and FLCs. The worst recovery (using the point at which the recovery rates cross 50% as a reference point) is seen at the lowest sky level ($5\text{--}7.5e^-/\text{pix}$), consistent with the most severely degraded SNR observed in Figure 5. Meanwhile the best recovery is typically at backgrounds between 20 and $50e^-$, consistent with the peak in SNR observed in Figure 5. At higher sky levels, the recovery begins to decline once more, but not reaching the worst levels again, at least in the range of backgrounds incorporated into this analysis.

To more clearly see the completeness limits as a function of sky background, Figure 7 plots the brightness corresponding to 50% recovery versus background level. Mirroring the trends in SNR versus background from Figure 5, we observe a rapid decline followed by a gradual increase in the 50% completeness level. A practical lesson from this analysis is the desirability that observations have background levels above $\sim 30e^-/\text{pix}$, but the SNR declines only modestly as backgrounds go above this level. However, users should be aware that the recovery of sources in the FLT/FLC images does not match that of the perfect CTE images until the backgrounds are greater than at least $\sim 100e^-/\text{pix}$.

Another notable feature is the worse recovery of sources in the FLC relative to the FLT at backgrounds below $\sim 50e^-/\text{pix}$, despite the fact that the FLC is intended to correct for CTE losses. One possible cause of this behavior may be the known noise amplification in




Figure 5: SNR measured within 2×2 pixel apertures for simulated point sources superimposed over an intrinsic sky background of $\sim 5e^-$, but with post-flash durations ranging from 0-14 seconds to increase the background. Different colored lines represent sources with different intrinsic brightness, f (in electrons), with the top and bottom panels focusing on fainter and brighter sources, respectively. The legend indicates the number of stars, n , that contribute to each colored line. Thick shaded bands show the measurements from the perfect-CTE images, overlaid for direct comparison with the results from the more realistic FLT and FLC mocks. The data shown are limited to the quadrants of each WFC CCD furthest from the serial registers.

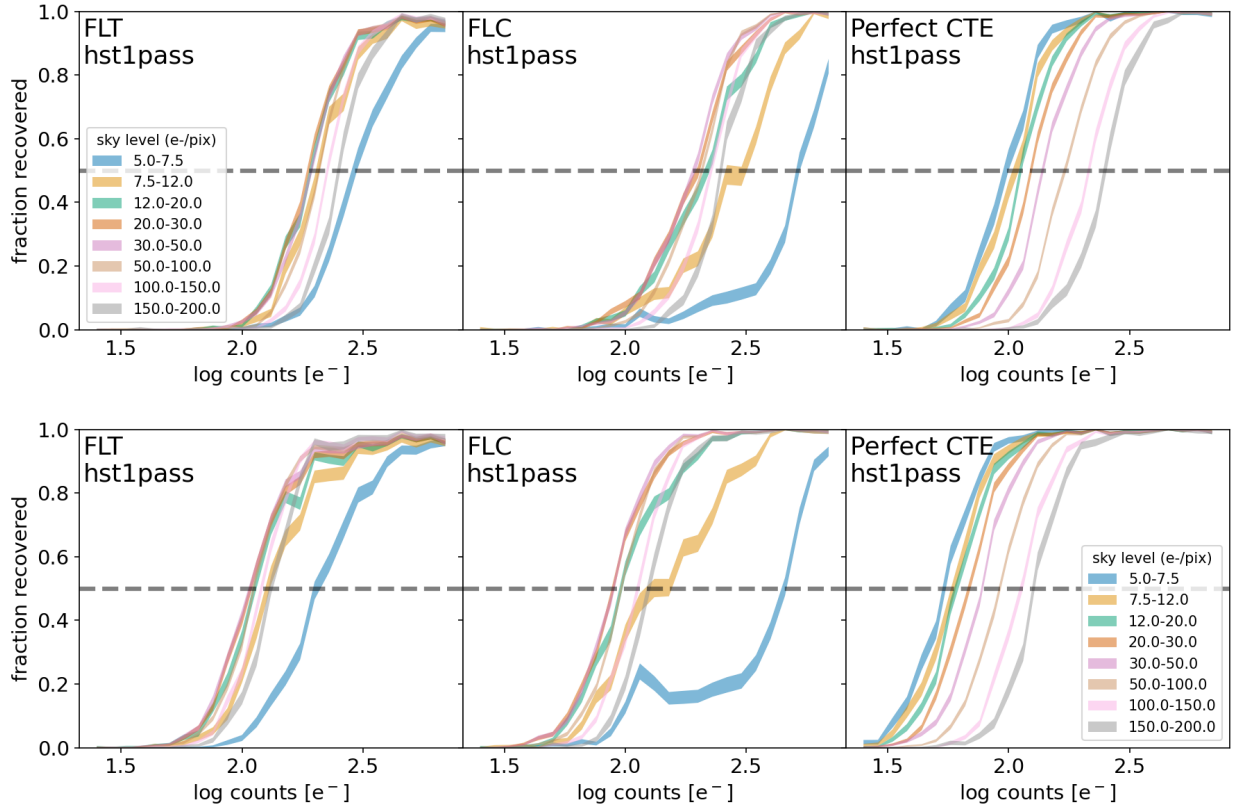


Figure 6: The fraction of recovered point sources using `hst1pass` in mock ACS/WFC single exposures (top) and 4-exposure stacks (bottom) as a function of source brightness (measured in 2×2 pixel apertures) and local background level. The analysis is limited to the quadrant of each ACS/WFC CCD furthest from the CCD serial registers. The dashed horizontal lines indicates where 50% of sources are recovered.

FLCs caused by the pixel-based CTE correction. The higher noise level would impact the detectability of sources by increasing overall noise level, but there may also be an increase in false detections (see Section 4) which could “crowd-out” genuine weak detections due to our isolation criteria of 5 pixels.

3.3.1 Performance on stacked images

Typical observations with ACS/WFC involve taking multiple dithered exposures to increase depth, mitigate the impact of cosmic rays and bad pixels, and improve PSF sampling. To consider the benefit of combining multiple exposures on source recovery, Figures 6 and 7 display the results for stacked images created by median-combining four exposures of the exact same field. These stacks crudely approximate the more typical (and recommended) 4-point dither strategy for ACS (the stacks used in our analysis are more akin to CR-SPLIT observations). However, we do not expect the lack of dithering to impact our key result. Most dithers in typical WFC observing programs are small ($<1\%$ of the field of view) so the CTE characteristics of all pixels containing a given source are highly similar anyways. However, it should be noted that our approach does not test whether improved PSF phase sampling gained from dithering has an impact on source detectability. Simply median combining the




Figure 7: The 50% recovery level (based on the data in Fig 6) for single exposures (left) and 4-exposure stacks (right) as a function of background level using `hst1pass`. Similar to the relationship between SNR and sky background, the best source recovery occurs when the background level is $\sim 30e^-$ per pixel.

exposures enables more direct comparison to the analysis of single exposures, and similarly, the intrinsic brightnesses and background levels shown in Figures 6 and 7 are the values from each single exposure, not summed over all exposures.

As expected, given the increased total observing time, the sensitivity of the stacks are improved compared to single exposures. The sensitivity of the FLC relative to FLT images improves more rapidly with increasing background, overtaking that of the FLTs at a much lower background level than for a single exposure. This is likely due to the image stacking reducing the influence of randomly located amplified noise from individual FLC exposures. The stacked FLCs also more rapidly converge with the sensitivity of an equivalent perfect-CTE image stack compared to individual exposures. However, because CTE impacts individual exposures, not the stacked result, the recovery of sources at low background levels is still extremely poor.

3.3.2 Source recovery with DAOSTarFinder

We next test source recovery using `DAOSTarFinder`. As noted in Section 2.3, this program is distinct from `hst1pass` in that it does not perform PSF fitting, and although we can filter results based on shape parameters (roundness and sharpness), at this stage we do not do so. As a result, one can consider the source detection rates and completeness limits shown below to be upper limits given that trimming the catalog of false detections could reject some real detections (see Section 4.2).

The recovered fractions of sources as a function of intrinsic brightness and background are shown in Figure 8. Trends are broadly similar to the analysis using `hst1pass`, but with a few exceptions. First, the maximum recovery rate with `DAOSTarFinder`, at least in the parameter space tested here, is lower than for `hst1pass`, maxing out at roughly 90% for FLTs (compared to $\sim 100\%$ when using `hst1pass`). The recovery rate may go higher for FLC and perfect-CTE images, but has not yet reached this level by $\sim 10^{2.75} = 500e^-$. As was the case with `hst1pass`, the recovery of sources in the FLCs is erratic, especially at low




Figure 8: Same as Figure 6 except using DAOSTarFinder for source detection.

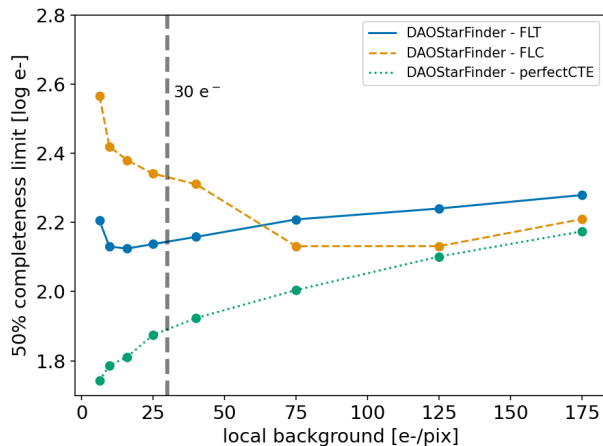


Figure 9: Same as Figure 7 but using DAOSTarFinder for source detection.

background levels. We again suspect noise amplification in FLCs to be responsible for this behavior.

The 50% completeness levels for single exposures as a function of background are shown in Figure 9. The trend for FLT shows the same decrease then increase in the completeness level that we saw in Figure 7, although the best completeness occurs at a lower sky background level than when using `hst1pass`. This difference may be due to the fact that we do not apply any shape restrictions on detections, whereas the analysis using `hst1pass` required $q < 0.5$ by default. In other words, the analysis here likely contains more marginal detections of real sources (as evident in the lower overall 50% completeness limit curves with DAOSTarFinder versus `hst1pass`), improving the detection rate at low backgrounds. The FLC also shows a decline in completeness level as a function of background, although much more gradual than for the FLT. It does increase again at high sky value, although we do not sample higher background levels so cannot confirm whether the trend persists. Like the analysis with `hst1pass` the completeness of sources in the FLT images is better than with FLC images at low sky background level, but in this case for backgrounds less than $75e^-/\text{pix}$.

4 False Detections

So far our analysis has focused on recovery of real sources, but false detections are another major consideration for blind searches in imaging data. In an ideal situation, one has an understanding of their completeness level for real sources and the contamination rate from false sources as a function of source brightness. The situation for ACS/WFC data is made more complicated by the fact that imperfect CTE in real images can create false “sources” that are inconsistent with simple Poisson noise. Similarly, the CTE correction to generate FLC images can amplify noise spikes, creating additional false detections. In the following section, we analyze the false detections from both `hst1pass` and `DAOStarFinder`, including their frequency and the parameter space they occupy. We will also discuss reasonable cuts users can apply to the catalogs to remove false detections that may arise. However, it should be noted that we do not conduct an exhaustive analysis of the best ways to trim the catalogs (or how to configure the source detection algorithms in the first place). Users’ needs differ, and some analyses may require high completeness at the expense of false detections, while others may need high purity at the expense of completeness. Our goal is to discuss what causes false detections and when users may need to be cautious of them.

4.1 `hst1pass`

A major benefit of `hst1pass` is its PSF fitting, which we find to be highly effective at removing false detections. There are essentially no false detections, across all background levels, when detecting point sources using `hst1pass` with the search parameters described above, i.e., setting `FMIN` to five times the average background noise level, and setting the maximum allowed goodness-of-fit parameter to the default $q_{max} = 0.5$ (Figure 10). We even recover a significant number of real sources with $SNR < 5$. Keeping `FMIN` the same but increasing q_{max} to 1 unsurprisingly introduces a large number of false detections, which mostly occupy a region defined by $q > 0.5$ and $SNR < 5$, regardless of whether we are considering FLT, FLC, or perfect-CTE images (Figure 10). It can be seen why $q_{max} = 0.5$ is a wise choice to create a clean sample, with the caveat that this cut rejects many real sources. Although one could ignore q and instead require e.g., $SNR > 5$, there are some high SNR false detections that would have been rejected by requiring $q < 0.5$, especially in FLTs.

Figure 11 shows the fraction of false detections as a function of background level when allowing q to be as high as 1.0. We find the number of false detections is roughly constant for FLT, FLC, and perfect-CTE images, and it also has very little dependence on sky background. False detections for FLC and perfect-CTE images also have the same very mild dependence on the number of y shifts for readout (Figure 12), while behavior for FLTs has a reversed behavior. The lack of any relationship between false detections and background level, the lack of any increase in false detections with increasing number of y shifts in FLTs, and the fact that the perfect-CTE exposures show false detection rates comparable to FLT and FLC images, all suggests the false detections in `hst1pass` are *not* associated with CTE. As we will see in the following section, the behavior for `DAOStarFinder` is notably different.

We find that adjusting the `FMIN` parameter in `hst1pass`, which sets the amplitude of pixels where `hst1pass` attempts to fit a PSF, has little impact on the number of reliably




Figure 10: Distribution of q vs. SNR for sources detected by `hst1pass`, with real and false detections in different colors. Contours are logarithmically spaced between a probability density of ~ 0.03 and 1.0.




Figure 11: (top) The number of total, real, and false detections using `hst1pass` with $q < 1$. (bottom) The rate of false detections. The vertical dashed line indicates the minimum recommended background of $30e^-/\text{pix}$.




Figure 12: Distribution of false sources as a function of y shifts when using `hst1pass` and allowing q to go as high as 1.0. There is no SNR cut on the data.




Figure 13: (top) Total, true, and false detections found by `DAOStarFinder` in FLT, FLC, and perfect-CTE mock images as a function of the local background level. (bottom) The false detection rate in each type of image. The vertical dashed line indicates the minimum recommended background of $30e^-/\text{pix}$. The dashed curves indicate the numbers/rates after removing false detections using the cuts illustrated in Figure 15.

detected sources. When lowering `FMIN` to 3 times the image noise level, we recovered only 1.7% more real sources all of which had $q > 0.5$, but we also doubled the false detections with $q > 0.5$, which vastly overwhelm any real sources detected there. While one can try to identify sources deeper into the noise, it is clearly not the ideal use-case for `hst1pass` as it may be extremely difficult to separate them from the high number of false detections.

4.2 DAOStarFinder

`DAOStarFinder` does not employ PSF fitting, and although it can incorporate constraints based on source shape, we do not apply these upon execution. For this reason, the list of sources returned by `DAOStarFinder` is heavily contaminated with false detections, especially in FLC and FLT images (Figure 13). The rate of false detections is globally high at all sky levels in FLTs, and gradually declines with increasing sky value. The false detection rate in FLCs is extremely high at low sky levels, but decreases rapidly as sky level increases. Meanwhile, the perfect-CTE images contain very few false detections in comparison. This sudden drop suggests the false detections in FLTs and FLCs are due to CTE trails, or in the case of FLCs, may have a contribution from noise amplification due to the CTE-correction algorithm. The link between false detections and CTE is evident in the dependence between the false detection rate and number of y transfers, while there is essentially no dependence for the perfect-CTE images (Figure 14). Additionally, the distribution of x positions for false detections in FLT images is spiky with typical separations consistent with the x separation of our mock stars, again consistent with false detections arising from trails behind real sources.

In reality, the majority of these false detections will be removed through basic selections on the catalog prior to any analysis. As noted above, the exact cuts will depend on the science goals of the analysis, but here we show simple cuts that avoid parameter space where false detections overwhelm real detections. We find that true and false detections are best

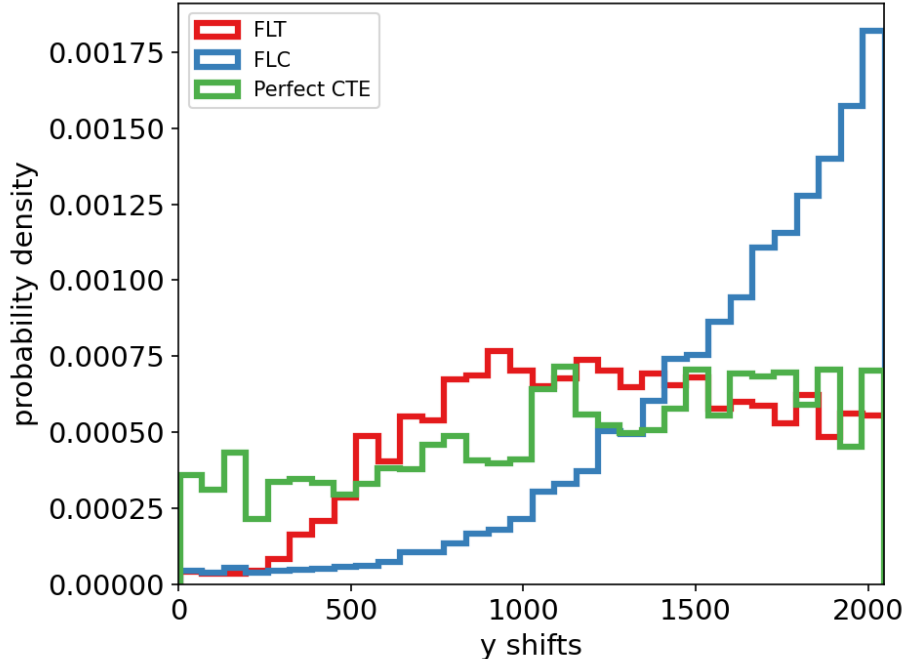


Figure 14: Distribution of false sources from DA0StarFinder as a function of y shifts.

separated based on their sharpness and SNR (measured within a 2×2 pixel aperture), but as shown in Figure 15, the exact divisions depend on the type of image. For FLTs, the true and false sources create two well-defined loci that can be approximately separated by the linear equation $s = -0.084SNR + 1.04$, where s is sharpness. For FLCs, most false sources can be rejected by requiring $s > 0.7$ and $SNR > 5$. The distribution of true and false sources in the perfect-CTE data is qualitatively similar to FLCs, but there are significantly fewer false detections. A similar cut as used for the FLCs would likely be appropriate.

Applying these selections is generally very successful at removing false detections, as shown in the dashed curves of Fig 13. After these cuts, the false detection rate is 5-8% for FLTs, <5% for FLCs (except at the very lowest sky levels), and < 0.3% for the perfect-CTE images.

5 Conclusions and recommendations for observers

By using a pipeline to generate mock ACS/WFC observations, we explore how imperfect CTE impacts the SNR and ability to detect point sources in ACS/WFC imaging data using two different programs, `hst1pass` and `DA0StarFinder`. We also explore what situations lead to false detections in each case. Our conclusions are as follows:

- For sources less than a few hundred counts in brightness, the SNR in a single exposure is maximized around a background level of $\sim 30e^-$. Below this level, SNR drops severely due to CTE losses, while above this value, SNR drops gradually due to the increasing contribution of background Poisson noise. Brighter sources do not show a strong increase in SNR with increasing background at low levels.




Figure 15: Distribution of sharpness and SNR for sources detected by DAOSTarFinder, with real and false detections in different colors. Dashed black lines show possible selections to eliminate the majority of false detections. Contours are the same as in Figure 10.

- Consistent with trends for SNR, source recovery is also maximized around a background level of $\sim 30e^-$.
- Despite the application of the pixel-based CTE correction, the SNR and source recovery in FLC files are not consistent with perfect CTE images until background levels are $> 100e^-/\text{pix}$.
- False detections (mostly at $\text{SNR} < 5$) will arise using `hst1pass` if the maximum goodness of fit parameter, q , is allowed to go above 0.5. These false detections do not generally appear associated with CTE.
- CTE effects drive many of the false detections found using `DAOStarFinder`, but most can be rejected using simple cuts on source SNR and sharpness.

Based on our results, we provide the following recommendations and reminders to users (some of which are reiterations of existing best practices):

- Be aware that the ACS ETC **does not** account for flux loss due to CTE.
- Ensure background levels will be above $30 e^-/\text{pixel}$, especially in regions of the CCDs far from the serial registers. Users can consult Anand et al. (2022) for recent information about ACS/WFC backgrounds in different filters. If obtaining background levels $> 30 e^-/\text{pixel}$ is difficult (e.g., for narrow-band observations), use post-flash to increase backgrounds to this level. Although adding post-flash adds an extra complication to data reduction and analysis (e.g., a non-uniform noise contribution across the CCDs), the improvements in SNR are significant.
- At low sky level, the pixel-based CTE correction does not recover all lost flux. In these cases, using FLTs and applying the photometric CTE correction from (Chiaberge & Ryon, 2022) may be the best path forward. However, be aware that this correction has not been calibrated for stars with brightness $< 250e^-$.
- Users should be cautious of the significant number of CTE-trailing induced false sources that can arise at low SNR thresholds. However simple selections on output catalogs should be able to yield clean samples, at least for the two detection algorithms used in this study.

Readers should keep in mind that our analysis of source recovery is subject to the exact configuration of `hst1pass` and `DAOStarFinder`. We have made only a few modifications from default settings (see Section 2.3), and alternative parameters may impact completeness limits, and/or purity of the output catalogs. Additionally, our analysis has only focused on point sources, these being the simplest targets observed and a natural starting point. However, much work with HST is focused on extended sources, and future work will explore whether CTE significantly impacts measurements of extended sources in ACS/WFC imaging data. Users should also keep in mind that our stars are centered on individual pixels, so we have not explored the impact of different pixel phase sampling. In reality, a star that is not centered on a pixel may have a broader, shallower profile, with different numbers of electrons in individual pixels (see e.g., Anderson & King 2000). The difference in sampling

will impact CTE losses and general detectability of that star in a manner distinct from if it is centered on a pixel. The impact of different pixel phase sampling will also be the subject of future investigation.

While this report has highlighted the negative impact of CTE on ACS/WFC data, high quality data can still be obtained with proper planning. Users are encouraged to contact the HST Help Desk⁶ with any specific questions regarding the impact of CTE on their existing or planned observations.

Acknowledgements

We thank Jay Anderson for his advise with using `hst1pass` and effective ACS/WFC PSFs. We also thank Gagandeep Anand, Nimish Hathi, Roberto Avila, Amy Jones, and Yotam Cohen for feedback on our analysis and this report. This work made use the Numpy (Harris et al., 2020), Astropy (Astropy Collaboration et al., 2013, 2018, 2022), Matplotlib (Hunter, 2007), photutils (Bradley et al., 2023), and scipy (Virtanen et al., 2020) Python packages.

References

- Anand, G., Grogin, N., & Anderson, J. 2022, Revisiting ACS/WFC Sky Backgrounds, Instrument Science Report ACS 2022-1, 30 pages
- Anderson, J. 2016, Empirical Models for the WFC3/IR PSF, Instrument Science Report WFC3 2016-12, 42 pages
- . 2022, One-Pass HST Photometry with `hst1pass`, Instrument Science Report ACS 2022-02
- Anderson, J., & Bedin, L. R. 2010, , 122, 1035
- Anderson, J., & King, I. R. 2000, , 112, 1360
- Anderson, J., & Ryon, J. E. 2018, Improving the Pixel-Based CTE-correction Model for ACS/WFC, Instrument Science Report ACS 2018-04, 37 pages
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, , 558, A33
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, , 156, 123
- Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *apj*, 935, 167
- Bradley, L., Sipőcz, B., Robitaille, T., et al. 2023, *astropy/photutils*: 1.8.0, doi:10.5281/zenodo.7946442
- Chiaberge, M. 2012, A new accurate CTE photometric correction formula for ACS/WFC, Instrument Science Report ACS 2012-05, 23 pages

⁶<https://stsci.service-now.com/hst>

- Chiaberge, M., & Ryon, J. 2022, ACS/WFC CTE photometric correction: improved model for bright point sources, Instrument Science Report ACS 2022-6, 20 pages
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357
- Hunter, J. D. 2007, Computing in Science & Engineering, 9, 90
- Miles, N. 2018, Updates to Post-Flash Calibration for the Advanced Camera for Surveys Wide Field Channel, Instrument Science Report ACS 2018-02, 22 pages
- Ogaz, S., Chiaberge, M., & Grogin, N. A. 2014, Post-Flash Capabilities of the Advanced Camera for Surveys Wide Field Channel (ACS/WFC), Instrument Science Report ACS 2014-01, 11 pages
- Ryon, J. E., & Stark, D. V. 2023, in ACS Instrument Handbook for Cycle 32 v. 23.0, Vol. 23, 23
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261
- Williams, R. E., Blacker, B., Dickinson, M., et al. 1996, , 112, 1335