



STScI | SPACE TELESCOPE
SCIENCE INSTITUTE

Instrument Science Report COS 2025-10(v1)

Determining X-Walk Corrections for the COS FUV Detector

Sten Hasselquist¹, Thomas Ake¹, Dzhuliya Dashtamirova¹, John Debes^{1,2}, Gisella De Rosa¹, David French¹, Darshan Kakkad¹, Nick Indriolo^{1,2}, Rachel Plesha¹, David Sahnou¹

¹ Space Telescope Science Institute, Baltimore, MD

² AURA for the European Space Agency

11 July 2025

ABSTRACT

The COS FUV detector comprises two microchannel plate segments (FUVA, FUVB) with cross delay line anodes to report the dispersion (x) and cross-dispersion (y) location of the charge cloud generated from a photon event. The electronics used to measure the locations were optimized during ground testing, but as the detector ages with usage, these positions shift due to the lowered pulse height amplitude (PHA) of the charge cloud, a phenomenon known as “walk”. In this work, we determine the corrections needed to remove walk in the x direction. We find that, while the shifts are ~ 1 -2 pixels for photon events at and near the gain nominal values of $PHA = 10$ -12, events with $PHA < 7$ can have X-walk of up to ~ 7 pixels, particularly low PHA events occurring on the left side of FUVA. We have used these measurements to create a new XWLKFILE reference file that contains the corrections to be applied by the CalCOS XWLKCORR module to correct these shifts.

Contents

1. Introduction	2
2. Data Used and Calibration Procedure	3
2.1 Mock Dataset	6
3. Measurements	6
3.1 Reference Spectra	6
3.2 OSM1 Drift	7
3.3 Identifying Measurement Windows	8
3.4 Measurement Uncertainties	8
3.5 Edge Measurements	9
4. Measurement Results	11
4.1 Testing on Mock Dataset	11
4.2 Program 12793 Measurements	11
5. Construction of the XWLKFILE	13
6. Summary	18
Acknowledgements	18
Change History for COS ISR 2025-10	20
References	20

1. Introduction

The Cosmic Origins Spectrograph (COS) far-ultraviolet (FUV) detector is a photon-counting device utilizing two sets of microchannel plate (MCP) stacks to register incident photon events from the FUV spectrograph gratings. It employs cross delay line readouts to determine the dispersion (x) and cross-dispersion (y) position of where a photon is incident upon the MCPs (Indriolo 2025). Each segment (FUVA and FUVB) is digitized such that events are assigned pixel locations within a 16384×1024 (x × y) grid. However, the accuracy to which these positions can be measured by the detector electronics depends on the size, or pulse height amplitude (PHA), of the charge cloud generated by the MCPs. Charge clouds with lower or higher PHAs than those calibrated during ground testing have their positions systematically miscalculated in both the x and y directions, shifts referred to as “walk”. Specifically, because the number of electrons generated by the MCPs, i.e, the gain, decreases with usage in different regions due to charge depletion from repeated illumination, it is important to determine the variation of walk with PHA level so that the positions can be corrected.

While previous work has been done to measure walk in the y direction (Sahnou et al. 2011), in this work we determine the walk in x and produce corrections that

will improve the wavelength accuracy for FUV spectra regardless of gain level. To measure these shifts, dx , we use data from a program (PID 12793) that was designed to determine the nominal operating high voltage levels at lifetime position 2 (LP2), but also included lamp exposures at each HV for the purpose of measuring geometric and walk corrections. Sahnou et al. (2018) demonstrated that X-walk can indeed be measured from this dataset. We expand on the Sahnou et al. (2018) analysis and measure dx for a wide range of detector locations and PHAs for both FUVB and FUVB, interpolating and extrapolating these measurements to create an XWLKFILE reference file that is used in the calibration pipeline, CalCOS, to correct these X-walk shifts. In a companion ISR (Hasselquist et al. 2025), we describe measuring and applying an updated correction in the cross-dispersion direction (Y-walk).

Throughout this work we assume that X-walk is not dependent on the cross-dispersion location on the detector, so that the X-walk correction measured at one cross-dispersion location (in this case the LP2 primary science aperture region) is applicable to all other cross-dispersion locations. While this assumption may not be entirely correct, because the travel time of a pulse in the cross dispersion direction along the delay-line anode is small relative to the along-dispersion direction when reporting the x position, any additional shifts coming from different cross-dispersion locations should be smaller. For FUVB, we do not observe an obvious y-dependence on X-walk for gain-sagged Lyman α regions of the detector at LPs 1–6. One could also analyze hot spots across the COS detector to further test this assumption, as these regions have a temporary drop in gain because the count rate is higher than the rate at which charge can be replenished in the MCP pores, but such an analysis is beyond the scope of our present work.

The data and calibration procedure used for our analysis are discussed in Section 2. The measurement methods are presented in Section 3. Measurement results are presented in Section 4, where we first show the outcome of the code applied to a mock dataset in Section 4.1 followed by the results from the Program 12793 data in Section 4.2. The adopted interpolation and extrapolation methods used to create an XWLKFILE from the measurements are discussed in Section 5. We summarize the final results in Section 6, and provide an example of the improvements on the COS data. A detailed discussion of the improvements can be found in French et al. (2025).

2. Data Used and Calibration Procedure

To measure X-walk we used data from Program 12793, which was a program designed to determine the detector high voltage (HV) levels for LP2. We have used Pt-Ne wavelength calibration lamp exposures taken at a range of commanded HV levels, from 143 to 178 on FUVB and 142 to 175 on FUVB, using grating G160M, cenwave 1600, and FP-POS=3. This means that high S/N wavecal spectra are available for a wide range of PHAs at a fixed location on the detector, providing a useful dataset to measure X-walk. As described in detail in Section 3.1, we use only the visits (shown in Table 1) that contain exposures at the highest HVs to mitigate measuring X-walk from reference spectra that might still be affected by X-walk.

Table 1. Table adapted from COS TIR 2018-03 that shows the structure of each visit used to measure X-walk. The HVA/HVB values from the exposures used to create the reference spectra (Section 3.1) are bolded.

Visit	Exposures	Date	Exposure time (s)	FP-POS	APERYPOS	HVA	HVB
01	6	2/27/12	220	3	255.1	143,147,152,157,169, 175	142,147,153,162, 168,172
10	6	8/27/12	220	3	255.1	143,147,152,157,169, 175	142,144,153,162, 168,172
11	6	8/28/12	220	3	255.1	157,167,169,172, 175,178	153,163,168,169, 172,175
12	6	8/28/12	220	3	191.1	147,157,167,172, 175,178	144,153,163,169, 172,175

For any COS exposure, each pixel has a distribution of PHA values from the photon events, with most having a PHA close to the peak of the distribution, which is taken to be a measure of the gain there (often referred to as the “modal gain”). For our analysis, the exposures are initially calibrated with the standard switches for LP2 data calibration, including the default CalCOS PHA filtering set to include PHA = 2 to 23, as determined by the lower limit threshold (“LLT”) and upper limit threshold (“ULT”) data columns of the PHATAB reference file (see e.g., [Section 3.7.7 of the COS Data Handbook](#)). The x1d files of the highest HV exposures are combined to make reference spectra. Then, the data are reprocessed 32 times using a series of 32 custom PHATABs with LLT and ULT values chosen such that the data are calibrated using only events from a single PHA, from PHA = 0 to PHA = 31. This results in 32 individual PHA-separated spectra for each exposure, although in practice, any given exposure results in only a small handful of PHA-separated spectra with high enough S/N to measure the X-walk. In Figure 1 we show examples of some individual PHA-separated spectra for FUV A (top) and FUV B (bottom). The PHA = 1 and PHA = 2 spectra are shifted by several pixels in XCORR from the reference spectra (described in detail in Section 3.1) in each panel. The higher PHA spectra are closer in XCORR location to the reference spectra, which at these XCORR locations have a modal gain of ~ 13 .

Because the X-walk measurements are performed on extracted spectra, the events have already been corrected for geometric distortion. However, the geometric correction, which makes use of the GEOFILE reference file (see e.g., [Section 3.7.3 of the COS Data Handbook](#)), can be improved using X-walk corrected data. Therefore, there is an iteration process we follow, where we measure an X-walk correction from data calibrated with the existing GEOFILE, use X-walk corrected data to derive a new GEOFILE, then use this new GEOFILE to re-derive an X-walk correction. Because our X-walk measurements changed little except for at the extreme edges of each segment, we found that a single iteration was all that was required. The overall process for iterating between creating GEOFILES and X/YWLKFILES is described in more detail in Indriolo et al. (2025).

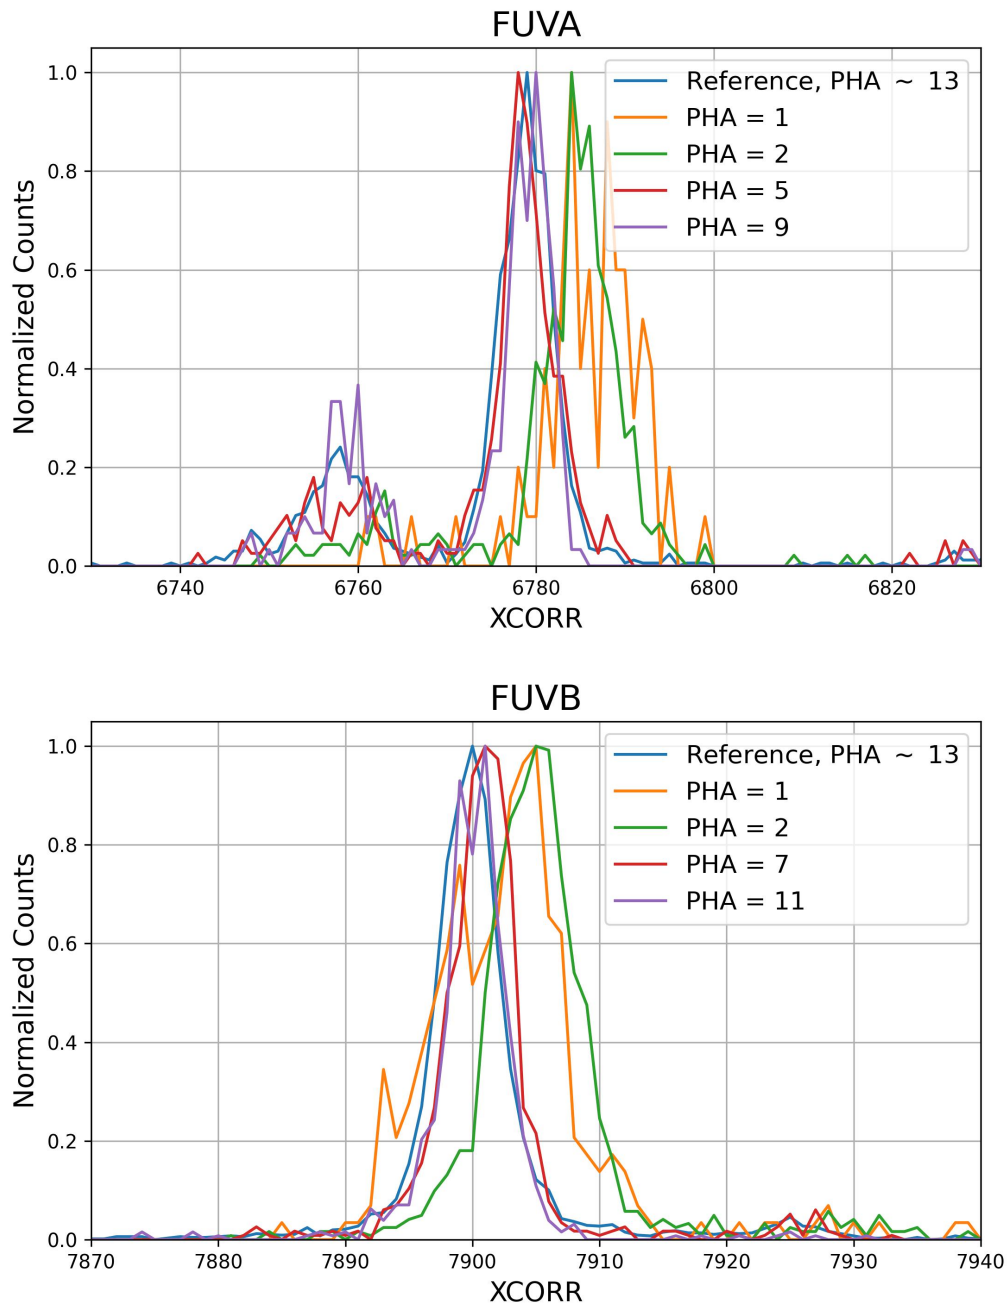


Figure 1. PHA-separated spectra compared to a reference spectrum for Visit 01 of Program 12793 for FUVB (top) and FUVB (bottom). Each spectrum is normalized by the maximum counts in this XCORR region. The lower PHA spectra are shifted in XCORR relative to the higher PHA spectra, demonstrating the effect of X-walk.

2.1 Mock Dataset

We also make use of a “mock” dataset to verify that the measurement code described in Section 3 is able to accurately measure a known input walk, and to help understand if the uncertainties we derive are a correct estimation of the precision of the measurements. To do this, we first generate a fake spectrum that consists of ~ 100 spectral features that are randomly spaced 5–150 pixels apart. Each feature has a Gaussian line profile with a standard deviation randomly chosen between 1.5 and 3.5 pixels and an intensity scaling factor that varies by a factor of 10. From this spectrum we then generate an events list by creating a number of photon events at each spectral feature that has a distribution of PHAs as set by the modal gain. For simplicity, we do not consider modal gain variations across the detector. For each Program 12793 exposure, we generate a mock events list with the modal gain set to the global modal gain of the data. We then apply a known X-walk to the events list, shifting the x location of each event based on its PHA. From each mock events list, a 1D spectrum is created for all PHA as well as for individual PHAs, with the NET counts of the mock spectra scaled such the S/N of the spectrum extracted across all PHAs matches the S/N of the corresponding Program 12793 exposure.

3. Measurements

Once the data have undergone the custom calibration discussed above (both the Program 12793 data and the mock dataset), the process for measuring X-walk is as follows:

1. Create a reference spectrum for each visit (Section 3.1).
2. Account for drift of the Optics Select Mechanism 1 (OSM1) during each visit (Section 3.2).
3. Identify individual emission-line windows to allow for a measurement of X-walk at different x locations on the detector and measure the shift in the individual PHA-separated spectrum relative to the reference spectrum via cross-correlation at the emission-line windows (Section 3.3).
4. Estimate the uncertainty of each measurement by generating 5000 realizations of both PHA-separated spectra and reference spectra and repeating the cross-correlation measurement (Section 3.4).
5. Use the location of the apparent edges of the detector segments to measure X-walk at the edges of the active area (Section 3.5).

3.1 Reference Spectra

To measure X-walk, we need a spectrum that is assumed to be unaffected by it, and therefore serves as a reference from which the X-walk can be measured. Ideally this

spectrum would be of high S/N such that spectral features of various intensity are well-measured across the entire detector. The reference spectra for each visit are created by combining spectra from all exposures taken at $HV \geq 173$ for FUVB and at $HV \geq 168$ for FUVB. These thresholds were chosen to avoid creating reference spectra that were affected by X-walk. They were determined by deriving a dispersion solution from the highest HV exposure in a given visit, applying it to the lower HV exposures, computing the differences between the measured and dispersion-solution line positions in the lower HV exposures, and evaluating the HV at which large residuals begin to appear ($\sigma > 2$ pixels). A detailed version of this test is explained further in French et al. (2024).

3.2 OSM1 Drift

Before measuring the XCORR shift in the individual PHA spectra relative to a reference spectrum, we need to remove the shift due to OSM1 drift during the visit. For the Program 12793 data the OSM1 was not commanded to move throughout a visit, but the position of the mechanism can shift slightly after an initial move, causing a drift in the x direction. The drift magnitude is typically ~ 2 pixels within the first 200s of an exposure after a major OSM1 movement (James et al. 2023), but it can be as high as ~ 6 pixels after 600s (Rowlands et al. 2024). For Program 12793 the exposures were taken in increasing order of HV, such that the first exposure was taken at the lowest value. Thus, the exposure at lowest gain is expected to be maximally shifted relative to the last exposure due to the effects of both X-walk and OSM1 drift, with subsequent exposures decreasingly affected as the mechanism settles.

To measure the OSM1 drift for each exposure, we cross-correlate individual PHA spectra (like those shown in Figure 1) of each level across neighboring exposures at the same level, as any shift between consecutive individual PHA spectra should be due to OSM1 drift and not X-walk. Only adjacent exposures in time typically have overlapping individual PHA spectra for which this can be done since they have similar HV values and therefore similar gain. After cross-correlating, we “chain” the measurements together, adding up the shifts between consecutive exposures until we are able to determine the overall drift between the first exposure and the last exposure in a visit. Because we end up doing this for a few PHAs, the final shifts we use are the average across all PHAs and segment. The adopted OSM1 drift measurements for each exposure in each visit are shown in Table 2. The final exposure in each visit is not listed as this exposure is used as the reference spectrum, and thus has no shift.

Values of “NaN” in Table 2 indicate situations where there was insufficient overlap in the PHA-separated spectra to measure a reliable value. These generally only occurred for a few FUVB exposures where the HV changes from exposure to exposure were larger than other exposures (see Table 1). In the cases where a shift from FUVB could not be measured, only the shift measured from FUVB is used, as the shift should be the same for each segment. Because we are dealing with much noisier, individual PHA spectra, the cross-correlation signals are often noisier, and the difference between

Table 2. OSM1 drift measurements

Exposure	HVLEVELA/B	FUVA OSM1 Drift (pix.)	FUVB OSM1 Drift (pix.)
lhx301enq	143/142	-4.2	-2.0
lhx301eqq	147/147	-1.6	-0.5
lhx301etq	152/153	-1.2	NaN
lhx301exq	157/162	-0.7	0.0
lhx301ezq	169/168	0.1	0.2
lhx310clq	143/142	-3.0	-6.3
lhx310cq	147/144	-1.8	-5.4
lhx310cvq	152/153	-0.9	-2.5
lhx310d1q	157/162	-0.1	-0.7
lhx310d4q	169/168	0.1	-0.2
lhx311gnq	157/153	-2.8	NaN
lhx311grq	167/163	-0.8	-0.9
lhx311gxq	169/168	-0.2	-0.2
lhx311h0q	172/169	0.1	0.1
lhx311h8q	175/172	0.2	0.1
lhx312izq	147/144	-1.4	NaN
lhx312j1q	157/153	-1.5	NaN
lhx312j7q	167/163	0.1	-0.2
lhx312jbq	172/169	0.7	0.3
lhx312jjq	175/172	0.5	0.2

FUVA and FUVB measurements suggest $\sim 1\text{--}2$ pixel uncertainties in this OSM1 drift measurement method. However, we verified that this method successfully decouples our OSM1 drift measurement from walk by measuring walk from a mock dataset with a known walk and OSM1 drift applied.

3.3 Identifying Measurement Windows

Ultimately, the `numpy.correlate` function is used to measure the shift between a given PHA spectrum and its corresponding reference spectrum. This is done for individual XCORR windows to allow for walk measurements as a function of x . For each reference spectrum, `skimage.filters.apply_hysteresis_threshold` is run to identify spectral lines in windows of size 5 \AA . The cross-correlation is then done at each of these windows, allowing for a walk measurement as a function of XCORR. When we report the measurements we use the mean position of these windows as our x coordinate.

3.4 Measurement Uncertainties

To estimate the uncertainty for each measurement, we follow techniques similar to those used in Peterson et al. (1998) and Maoz & Netzer (1989) and generate 5000

realizations for each reference spectrum and PHA-separated spectrum, and repeat the cross-correlation measurement. Each spectrum is generated by randomly drawing from a normal distribution created at each pixel, with mean and standard deviation equal to the mean counts and mean uncertainty on the counts, respectively. This often results in asymmetric distributions, so each measurement is assigned an upper and lower error bar, corresponding to 84th percentile and 16th percentile, respectively.

These uncertainties are used in two ways. First, we use the mock dataset to help inform us what uncertainties result in erroneous measurements. We conservatively keep all measurements with $|\sigma_{84} - \sigma_{16}|/2 < 5$ pixels. Second, we use the uncertainties as weights in the spline interpolation, discussed in detail in Section 5.

3.5 Edge Measurements

Using this dataset, we can only measure detector walk where lines appear, so we are unable to measure shifts out to the edges of the active area. However, there is a physical edge to the detector, so there should be a boundary past which no counts should occur. Therefore, we can get additional walk measurements by determining where the apparent edge of the detector is for each PHA, and deriving the shift required to move the lower PHA values to the same edge location as the higher PHA measurements.

To do this, we first make large, combined corrtags from the following Program IDs: 12793, 12676, 16828, 16535, 13971, 16322, 15771, 15533, 14940, 14520, 14436, 13968, 13521, 13121, 12716, 12423, 11895, and the thermal vacuum testing data used to derive the geometric correction. These datasets contain exposures at a variety of commanded HV levels and gain sag, and were thus selected to include photon events at a wide range of PHAs across much of the detector. For each segment and each edge, we select a range in YCORR that is relatively straight in the cross-dispersion direction and free from large gain-sag holes that make measuring the edge at higher PHAs difficult. The YCORR ranges used for FUV A are 505–585 and 510–550 for the left and right edges, respectively, and for FUV B are 550–600 and 525–605. These were determined by eye. We then bin the counts in single pixel bins in XCORR and calculate a running derivative of size 5 pixels, subtracting the counts from the first and last pixels in each bin and using this as our measurement for the change in counts as a function of XCORR location. The left edges are determined to be where this derivative is maximum, and the right edge is where this derivative is minimum. An example of how the edge is determined for a few select PHAs is shown in Figure 2, with the right column demonstrating where the edge was found based on the derivatives of the middle column.

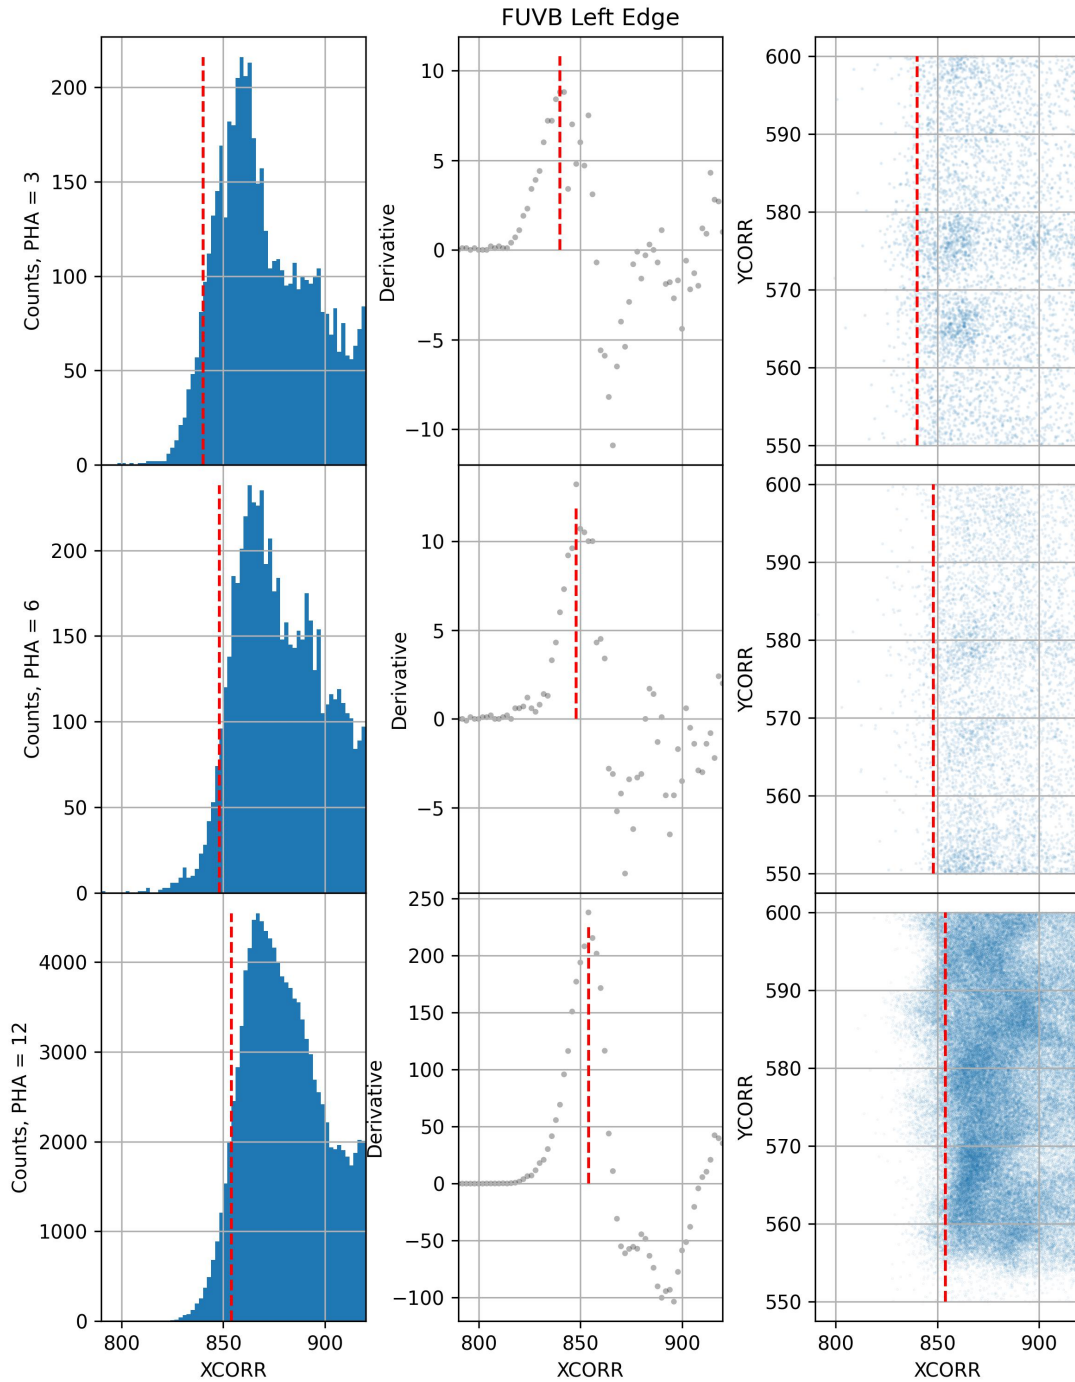


Figure 2. Schematic showing how the edge is measured for a few select PHAs (rows) at the left edge of FUVB. The left column shows the binned counts distribution for the edge. The middle column shows the running derivative calculated for this distribution. The right column shows a scatter plot of the counts. In each panel, a red dashed line marks the edge location, as determined by the maximum of the derivatives shown in the middle column. In this case, the apparent edge that we measure shifts to higher XCORR at higher PHA.

4. Measurement Results

We first run the measurement code on a mock dataset to ensure that the code is able to measure a known input walk, and understand to what precision it can be measured. These results are presented and discussed in Section 4.1. We then present and discuss the results from the main Program 12793 dataset in 4.2.

4.1 Testing on Mock Dataset

As described in more detail in Section 2.1, we first run the measurement code on a dataset that has a known walk applied so that we can analyze the accuracy with which the measurement code recovers the input walk. Results are shown in Figure 3. The top two panels of Figure 3 demonstrate that the measurement code qualitatively recovers the expected walk measurement at all XCORR locations and most PHAs. For FUVB, the code returns almost no measurements of walk at PHA = 0 and PHA = 1. Because we scale the flux of the mock dataset to match the real dataset, this means that, for FUVB, the emission lines are basically not present at the low PHA spectra. The FUVB dataset is a factor of ~ 4 higher in S/N, so some amount of walk is still measured at these lower PHAs. The FUVB mock dataset therefore ends up with a larger number of overall measurements, which is why some of the smaller-scale structure in the input walk is recovered better for FUVB than for FUVB.

The third row of Figure 3 shows that, for the PHAs that comprise the vast majority of the mock dataset ($1 < \text{PHA} < 16$), the code is able to return the expected measurement with a standard deviation of 0.42 pixels for FUVB and 0.45 pixels for FUVB. For FUVB, this is about what is expected based on the reported uncertainties, shown on the bottom row. However, for FUVB, the uncertainties appear to be underestimated. Given that the FUVB dataset is the same as FUVB but with higher S/N, it appears that there is some residual systematic uncertainty that causes the resulting measurements to be as good as ~ 0.45 pixels. Because we mostly use the uncertainties as a threshold for “bad” measurements and as relative weights for spline fitting, we do not further investigate these uncertainties. However, based on the FUVB panel of Figure 3, we do set a lower error threshold of 0.1 pixels for each measurement, as this uncertainty results in a better explanation of the 0.45 pixel standard deviation shown in the FUVB residual panel.

4.2 Program 12793 Measurements

Having verified that the measurement code is able to reliably measure the X-walk applied to a mock dataset that was generated to closely mimic the input dataset, we then ran it on the actual dataset. Results for all measurements across all visits are shown in Figure 4. The majority of measurements span the range $1 < \text{PHA} < 15$ for FUVB and $1 < \text{PHA} < 18$ for FUVB, with a higher density of measurements for FUVB than

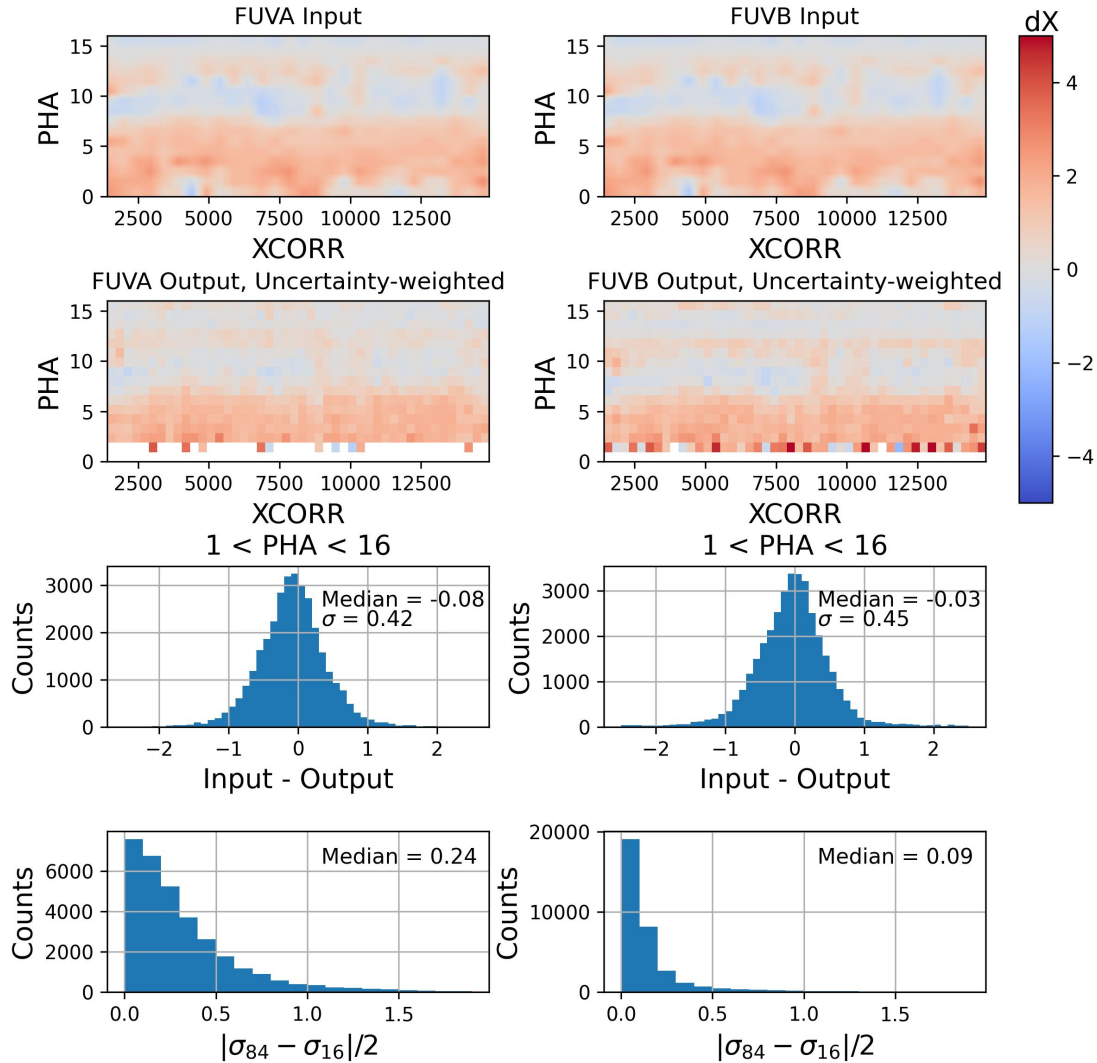


Figure 3. Measurement results from the mock dataset for FUVA (left) and FUVB (right). The FUVA and FUVB mock datasets are the same except that the FUVB mock dataset is generally a factor of ~ 4 higher in S/N , just like the real dataset. We remove any measurements with erroneously high walk measurements (> 100 pixels). The top row shows an $XCORR \times PHA$ image of the expected measurements based on the “anti-walk” applied to the mock dataset. The second row shows the weighted (by the measurement uncertainties) mean output measurements for a region that is 300 pixels in $XCORR$ at each PHA , with bins containing < 3 measurements appearing as white. The third row shows the distribution of residuals (input walk - output walk) for all measurements for $1 < PHA < 16$, with medians and standard deviations indicated in the upper-right of each panel. The bottom row shows the uncertainty distribution for these measurement, which we estimate as $|\sigma_{84} - \sigma_{16}|/2$.

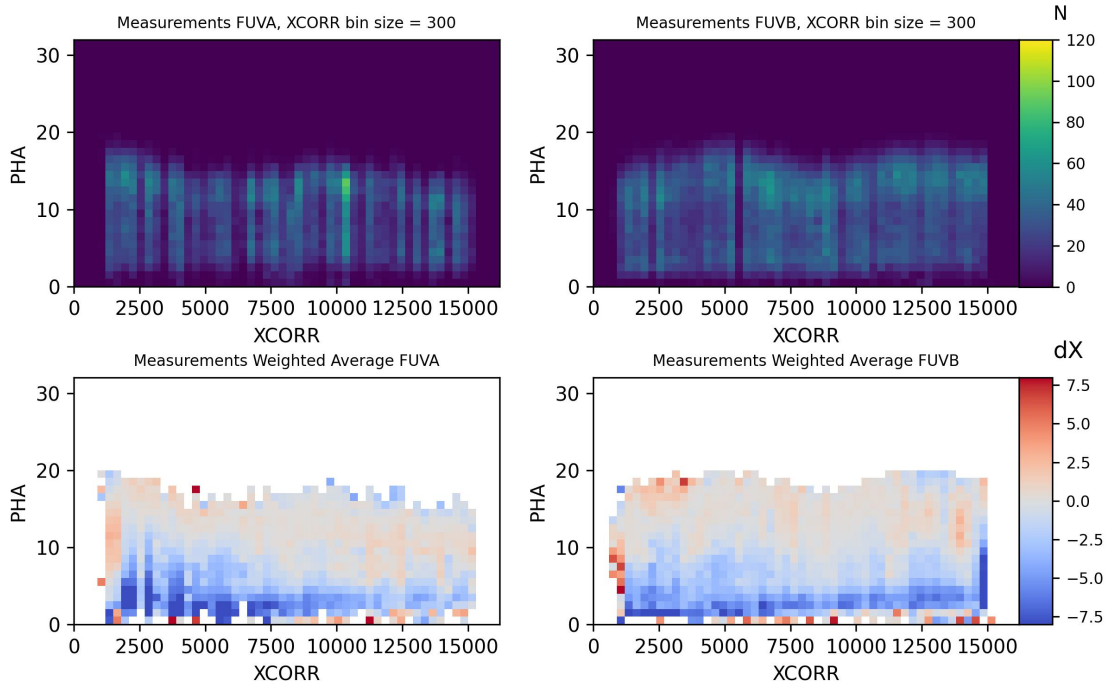


Figure 4. X-walk measurement results from Program 12793 data for FUVA (left) and FUVB (right). The top row shows the measurement density in bins of PHA and XCORR, where the bins are of size 1 in PHA and 300 pixels in XCORR. The bottom row shows the uncertainty-weighted average walk measurement of each of these bins.

for FUVA. As expected, the bottom row of Figure 4 shows that the X-walk is larger for lower PHA events, with shifts as high as 6–8 pixels for $\text{PHA} < 4$. For the higher PHAs, dx is generally much lower (1–2 pixels), except for at the edges.

5. Construction of the XWLKFILE

To apply X-walk corrections, the CalCOS calibration pipeline requires an XWLKFILE reference file of dimensions 16384×32 , such that each XCORR pixel and PHA has a correction (see [Section 3.4.7 of the COS Data Handbook](#)). Therefore, we must create an image from the dx measurements described in Section 4.2 and shown in Figure 4. The steps we take to achieve this transformation are:

1. Extrapolate dx measurements to higher PHAs that were not measured from the Program 12793 data.
2. Interpolate between the dx measurements, including interpolating between the Program 12793 measurements and the edge measurements described in Section 3.5.
3. Extrapolate dx values beyond the edges to the lowest and highest XCORR values.

4. Apply a zero point correction such that PHA = 11 for FUVB and PHA = 13 for FUVB have zero X-walk shift for all XCORR locations.

Each step is demonstrated in Figure 5, and described in detail below.

To extrapolate to higher PHAs that were not measured in the Program 12793 data, we first aggregate the measurements in bins of size 300 pixels in XCORR for each PHA, and compute the uncertainty-weighted mean shift in each bin. This is shown in the top row of Figure 5. For each XCORR location, we identify the highest PHA bin in which the number of measurements drops to 3 or fewer. We take the mean X-walk shift of the two PHA levels just below this bin, and set all PHAs from there upwards to PHA = 31 to that value, as demonstrated in the second row of Figure 5. We do not perform a linear extrapolation here, as the walk correction is predicted to be smaller with increasing PHA. Note that the vast majority of COS FUV science observations are obtained at PHAs below this extrapolated region, so we do not expect uncertainties in the extrapolation to have any impact on the quality of those calibrated data products.

To obtain X-walk values across the active area, we first perform a spline interpolation through all XCORR measurements at each PHA. We use a similar spline method as that used to interpolate in the GEOFILE (see e.g., Kakkad et al. 2025), using the `UnivariateSpline` function from the `scipy.interpolate` python package. We specify a fixed number of knots (15) by using `numpy.percentile` to identify the XCORR locations of evenly-spaced percentiles of measurements starting from 5% to 95%. A univariate spline is then fit to the knots, with their values determined by the inverse variance weighted means of the Program 12793 X-walk measurements. We determined the optimal number of knots to be 15 by fitting multiple splines to increasing numbers of knots from 3 to 50, calculating the standard deviation of the difference between resulting spline fit and the measurements. We chose 15 as the number of knots at which the standard deviations became marginally lower, i.e., increasing the number of knots beyond this does not result in a better fit. In the higher PHA regions (PHA > 18), where data are extrapolated only, we use 13 knots and the XCORR locations that were set as part of extrapolating to higher PHA values. An example of spline fits for some FUVB measurements are shown in Figure 6.

Once the walk values between the spline knots are determined, we then linearly interpolate the walk between the left-most and right-most knot locations and that of the edges measured in Section 3.5. Results of this interpolation are shown in the third row of Figure 5, which can be qualitatively compared to the binned measurements shown in the top row of Figure 5. We show the difference between these spline fits and the Program 12793 measurements in Figure 7, finding standard deviations (scaled from measuring the median absolute deviation) of 1.35 pixels for FUVB and 1.24 pixels for FUVB.

In principle, the BRFTAB, which defines the active area of each segment of the FUV detector, should be set such that the active area does not extend beyond the edges measured here. Still, we decided to extend the edge measurements to the lowest and highest XCORR values, as shown in the fourth row of Figure 5. Finally, since all of the measurements here are made with respect to an arbitrary reference spectrum, it is

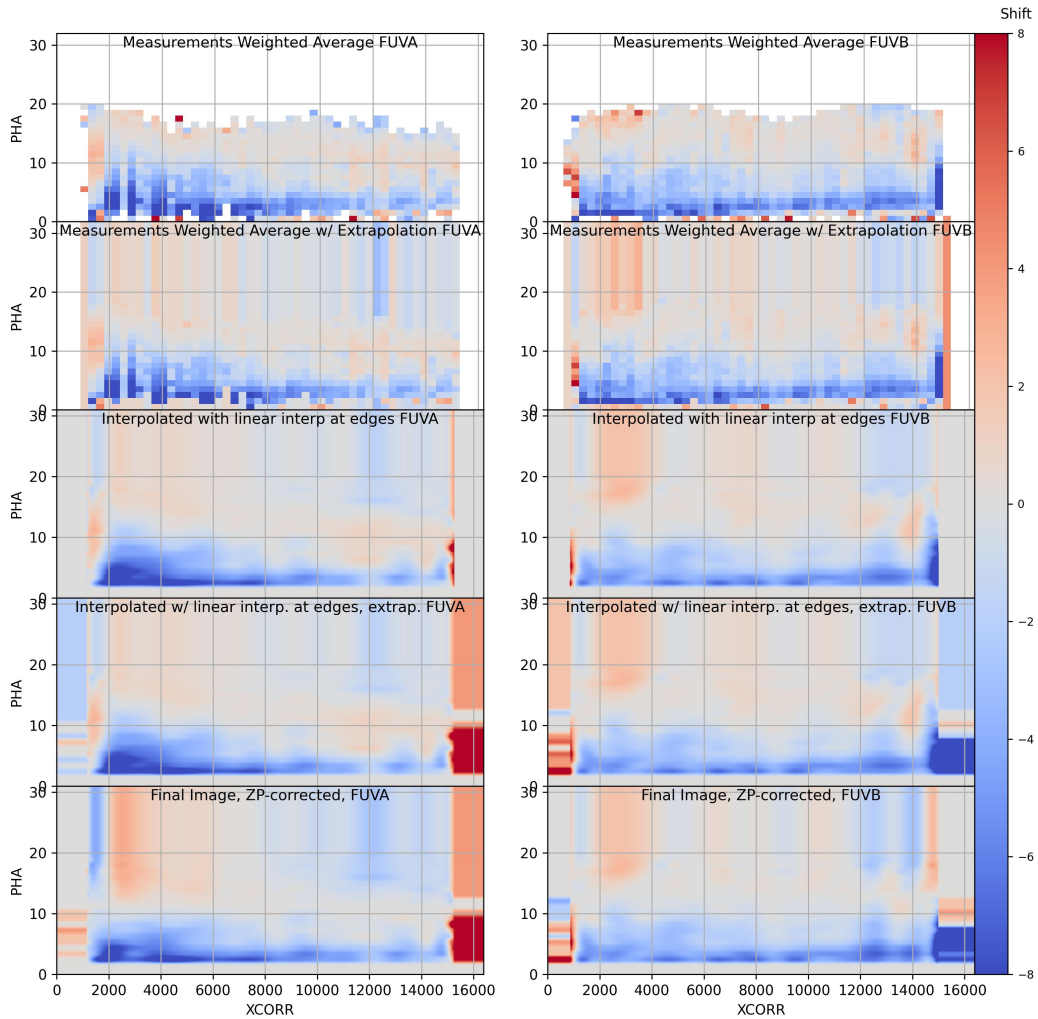


Figure 5. Extrapolation/interpolation steps to transform the measurements from the Program 12793 data to a reference file image of size 16384×32 for FUVB (left) and FUVB (right). The top row shows an image of the measurements, binned in size 1 in PHA and 300 pixels in XCORR, as was done in the bottom row of Figure 4. The second row demonstrates how measurements are extrapolated to high PHAs, where, for each XCORR bin, we take the mean of the two highest PHA bins with > 3 measurements and extrapolate that value to all higher PHAs. The third row shows the results of our 15 knot Univariate Spline interpolation, with a linear interpolation occurring between the first and final knots and their respective edge positions. The fourth row demonstrates how the edge measurements are extrapolated to the lowest and highest XCORR values beyond the edges of the active area. The final row shows the final reference files, where small zero-point normalizations have been made such that the walk for PHA = 11 and PHA = 13 for FUVB and FUVB, respectively, is zero across the segment.




Figure 6. X-walk measurements, dx and associated error bars (grey) with spline fits (red) for three PHAs for FUVB. The knot locations are shown as red circles. As described in detail in the text, we do not use the regions of the spline fits that diverge where they are unconstrained by data.

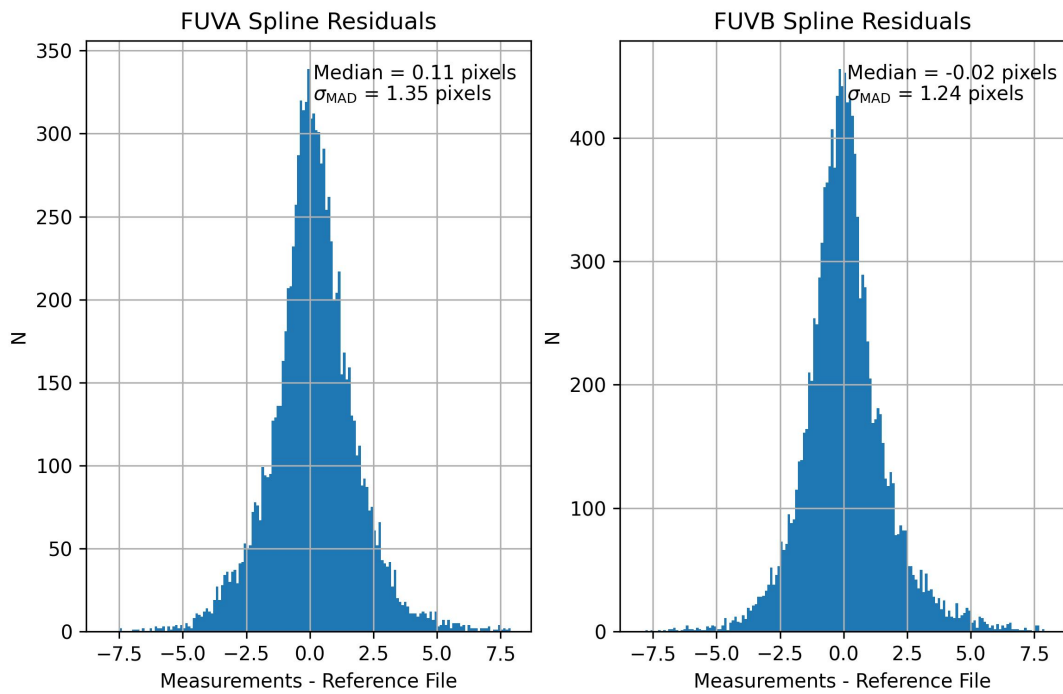


Figure 7. Differences between the spline fit and the X-walk measurements for measurements with $2 < \text{PHA} < 15$ and measurement uncertainties < 2 pixels for FUVA (left) and FUVB (right). Medians and standard deviations calculated from the median absolute deviation (σ_{MAD}) are indicated in the upper-right of each panel.

necessary to define a PHA level for which X-walk is defined to be zero at all XCORR. We choose PHA = 11 and PHA = 13 for FUVB and FUVB, respectively, and adjust dX values in the entire array as required to satisfy this definition. These PHA values were chosen because they are the modal gain values of the TV03 data set used for measuring the geometric distortion. The magnitude of these zero-point shifts are generally of order < 1 pixel, with the exceptions being the left edges ($XCORR < 3000$) of both FUVB and FUVB, where the magnitude of these shifts can be as high as 3 pixels for FUVB and 2 pixels for FUVB.

Due to the sparse coverage and relatively larger uncertainties, we do not use any X-walk measurements from PHA = 0 and PHA = 1 spectra, and instead set the X-walk shift for these PHAs to 0 at all XCORR. Note that these low PHA events are currently excluded from standard data products by the PHATAB, but any future attempts to recover these counts must take into account that their probable large X-walk shifts are not being accounted for in the XWLKFILE.

6. Summary

We have used Program 12793 data to measure X-walk at a wide range of XCORR locations and PHAs ($1 < PHA < 20$). Using a mock dataset, we verify that these measurements are generally accurate and precise to ~ 0.5 pixels, and that the method used to remove OSM1 drift is able to separate OSM1 drift from X-walk. We combine these measurements with X-walk determined at the segment edges and then interpolate/extrapolate to create an XWLKFILE reference file that can be used to correct COS data for X-walk. The largest X-walk corrections are found for PHA < 7 and are greatest at the left third of FUVB.

An example of how this XWLKFILE results in an improvement is by reprocessing the data used in the LP3 wavelength calibration, where Plesha et al. (2018) surmised that a systematic curvature in the wavelength residuals for FUVB was caused by gain sag. In Figure 8, the upper panel shows the difference between the measured position of emission lines and their expected position given the dispersion solution for data without a walk correction applied; the lower panel shows the results after walk is removed by the XWLKFILE. The standard deviation of the residuals decreases with the walk correction applied, and the overall curve found in the residuals is flattened. A more detailed description of the improvement can be found in French et al. (2024).

Acknowledgements

We thank Elaine Frazer for her initial work in writing and testing pieces of code that eventually were adapted and used in the walk measurement code. We thank Leonardo Dos Santos for careful review of an earlier draft of this ISR.

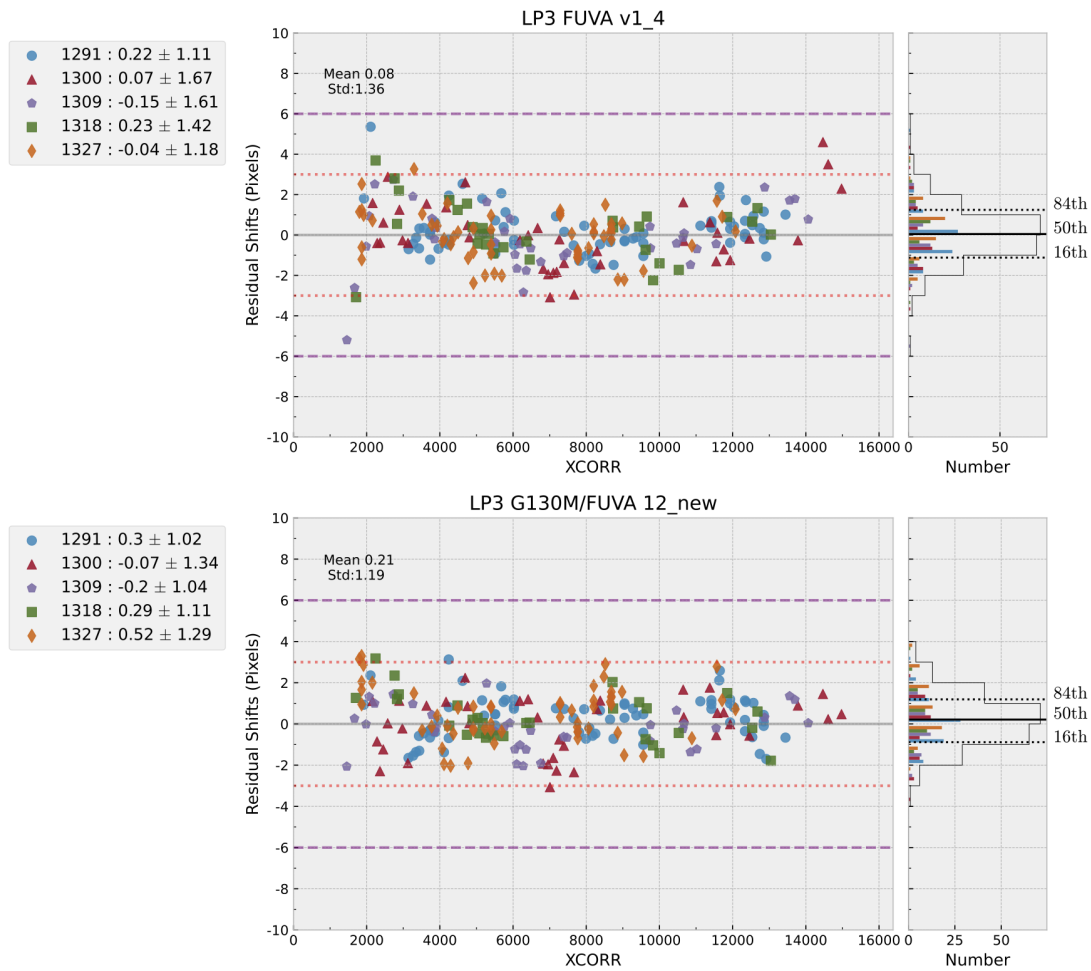


Figure 8. As described in detail in French et al. 2024, residuals of observed emission lines and their expected location from a dispersion solution as a function of XCORR for FUVA data without an X-walk correction applied (top) and data with an X-walk correction applied (bottom). The overall standard deviation of the residuals, as noted in the upper-left of each panel, decreases for the data with the X-walk applied. Moreover, the residual pattern is much flatter with XCORR for the data with the X-walk correction applied.

Change History for COS ISR 2025-10

Version 1: 11 July 2025- Original Document

References

- French, D., et al., 2025, COS Instrument Science Report 2025-XX
Hasselquist, S., et al., 2025, COS Instrument Science Report 2025-11
Indriolo, N., et al., 2025, COS Instrument Science Report 2025-07
James, B., 2023, COS Instrument Science Report 2023-10
Kakkad, D., et al. 2025, COS Instrument Science Report 2025-09
Maoz, D. & Netzer, H., 1989, MNRAS, Vol. 236, 21-29
Peterson, B. et al., 1998, PASP, Volume 110, Issue 748, pp. 660-670
Plesha, R. et al., COS Instrument Science Report 2018-24
Rowlands, K. et al., 2024, COS Instrument Science Report 2024-08
Sahnow, D., 2018, COS Technical Instrument Report 2018-03
Sahnow, D., et al., 2011, Proceedings of the SPIE, Volume 8145, id. 81450Q
Soderblom, D., et al. 2022, COS Data Handbook, Version 5.1, (Baltimore: STScI).