

Diagnostic measurements for the JWST field M4

JWST Galactic Centre pipeline — automated diagnostic write-up

2026-08-01

Abstract

This document characterises the astrometric, photometric and diffuse-background measurements delivered for M4 (JWST programme(s) 1979). It is an analysis of the finished data products, not a quality gate: it asks how well the measurements behave and where their limits are, rather than whether they pass a threshold. Every figure is generated from the products on disk by `jwst_gc_pipeline.diagnostics`, and every number quoted in the text is read back from the same computation that drew the figure.

1 Scope

Three related efforts characterise these data and it is worth saying at the outset which is which. The JWST-GC/data-qa repository checks *initial* data products, close to the telescope, and is oriented towards catching a bad exposure early. The astrometry paper is a technique-development document: it follows the iterative work of establishing how to measure a position at all in these fields. This write-up is neither. It takes the finished, released products for one field and measures their properties comprehensively, so that a user of the catalogue can see what precision they are entitled to assume.

Product gaps. No cross-band merged catalogue: colour-colour and cross-band astrometry figures are omitted. No per-filter catalogue for F150W2.

2 Products measured

The field directory is `/orange/adamginsburg/jwst/m4/`. The cross-band merge used here is `none` and the astrometric reference is `gaia_refcat.fits`.

Filter	Per-filter catalogue	Mosaic
F150W2	—	jw01979-o002.t001_nircam_clear-f150w2-merged_i2d.fits
F322W2	f322w2_nrcb_indivexp_merged_resbgsb_sub_m6_dao_basic.fits	jw01979-o002.t001_nircam_clear-f322w2-merged_i2d.fits

3 Overview

Figure 1 orients the rest of the document. The per-filter catalogues hold 25,847 measurements in total, the largest being F322W2 with 25,847 sources over roughly 31237 square arcseconds. The

peak surface density is 2.96 sources per square arcsecond and the typical occupied cell holds 0.76. At that density the mean separation between neighbours is comparable to the point-spread function, so crowding, not photon noise, sets the achievable precision over much of the field; the qfit and propagated-error diagnostics below should be read with that in mind. Number counts turn over at 2.11 in F322W2, which is the practical completeness limit of each band. These magnitudes are instrumental: this field has no cross-band merge from which to recover a photometric zero-point, so the scale is $-2.5 \log_{10}(\text{flux})$ in image units and only differences within a band are meaningful.




Figure 1: Overview of m4. Left: source surface density in 72×72 sky cells for F322W2, the filter with the most detections. Centre: the sky coverage of each filter, sub-sampled. Right: number counts per filter; the peak of each curve is the turnover magnitude, the practical completeness limit beyond which the counts are set by detectability rather than by the stellar population.

4 Astrometry

How the offsets were measured. Every offset in this section comes from offset-histogram stacking: all pairs within a search window contribute to a two-dimensional histogram of their separations and the peak is the offset. The alternative — matching each source to its nearest reference counterpart and taking the median — is not used and is prohibited in this pipeline. When the true shift exceeds the reference catalogue’s own nearest-neighbour spacing, nearest-neighbour matching pairs the wrong star, and the median of those wrong pairs collapses towards zero: the method reports agreement precisely when the frame is most badly misregistered. The search window is also swept rather than fixed, because a shift much larger than the window leaves no true pairs inside it and reads as noise rather than as a large offset.

4.1 Repeatability

Figure 2 shows the exposure-to-exposure scatter of each source about its merged position. The bright-end floor is $3.18''$ in F322W2. This floor is the single-exposure centroiding precision convolved with whatever residual frame-to-frame misregistration survived alignment; it is an upper limit on the latter, and the merged position of a star seen in N exposures is better than it by roughly $\sqrt{N}$. 19,552 measurements (76 per cent) report a scatter far below any achievable centroiding precision, in most cases identically zero. These are forced or seeded fits whose position was copied from a fixed seed and never re-measured per exposure, so the value is the absence of a measurement rather than a perfect one; they are excluded from the curves and from the floor above. At this

fraction the exclusion is not a detail: most of the catalogue carries no independent positional measurement, and the quoted floor describes the minority that does.

m4: internal astrometry




Figure 2: Internal astrometry of m4. Per filter: the scatter of the individual-exposure centroids about the merged position, against brightness; greyscale is source density, the coloured line and band are the running median and 16–84th percentile, and the dotted line marks the bright-decile floor quoted in the panel.

4.2 Absolute tie

Figure 3 ties the catalogue to its registered reference. Bulk offsets run from 140.5 to 140.5 m'' (largest: F322W2). Mapped per tile, the median tile offset across filters is 139.5 m'' , the 95th percentile 335.4 m'' , and the worst single tile 721.2 m'' in F322W2. More than five per cent of tiles exceed 50 m'' , so this is not an isolated bad tile: it is a spatially varying registration residual of the kind a field-averaged number washes out, and it is exactly what the per-tile map exists to catch.

The overlap regions of the affected bands should be inspected before those positions are used at the tens of milliarcsecond level. 20 of 100 tiles across all filters (20 per cent) reached a tie only after the search window was widened. Those tiles hold too few sources for a coherent peak at the nominal window, so their offsets are excluded from the numbers above and are marked rather than drawn in the figure; a high fraction means the map is sparse, not that the field is misregistered.

m4: absolute astrometry vs gaia_refcat.fits




Figure 3: Absolute astrometry of m4 against its registered reference catalogue. Each per-filter panel maps the offset-histogram tie on a 12×12 tile grid (arrows, exaggerated; colour is the per-tile peak contrast), with the whole-field bulk tie quoted in the corner. Grey crosses mark tiles whose tie was only found by widening the search window and whose offset is therefore not trustworthy. The final panel collects the bulk ties. Offsets are measured by offset-histogram stacking, which is immune to the nearest-neighbour pairing collapse that a median against a dense reference suffers; the per-tile map is shown because a bulk offset near zero does not by itself mean the field is registered.

5 Photometry

5.1 Precision and depth

Figure 4 shows the fractional flux uncertainty against brightness. The 5σ depth — where the median σ_F/F reaches 0.2 — is 1.51 mag in F322W2, on the instrumental scale: no cross-band merge exists for this field, so no photometric zero-point could be recovered and the numbers are comparable within this field only. The uncertainty propagated from the exposure-to-exposure scatter differs from the fitter’s formal covariance by a median factor of 1.00. The two agree closely, which means the noise model is adequate and the fits are not being perturbed by exposure-dependent crowding.

5.2 Fit quality

Figure 5 shows the normalised PSF-fit residual. Median qfit is 0.561 in F322W2; the fraction of sources above the vetting threshold is 71.9 per cent in F322W2. Where that fraction is large the

m4: photometric precision and depth




Figure 4: Photometric precision of m4. Greyscale is source density; the solid line is the running median of the fitter’s formal uncertainty σ_F/F and the dashed line the same quantity propagated from the per-exposure scatter. The dotted horizontal line is $\sigma_F/F = 0.2$, and where the median crosses it is quoted as the 5σ depth. A propagated-to-formal ratio above unity means the exposures disagree by more than the fit covariance predicts.

catalogue is dominated by sources whose profile the model does not reproduce, which in these fields is normally blending rather than a defective PSF: the qfit distribution rises towards the faint end where neighbours are unresolved, not towards the bright end where a bad PSF model would show. Saturation touches at most 3.2 per cent of a band (F322W2); the census panel breaks that down into the flagged, core-replaced, gate-rejected and clip-corrected populations, which is the accounting needed before trusting any bright-star photometry.




Figure 5: PSF-fit quality for m4. Per filter, the normalised fit residual qfit against brightness (greyscale: source density; line: running median; dotted: the $qfit = 0.2$ vetting threshold). The final panel is the fraction of each filter’s catalogue touched by the saturation-handling paths: flagged saturated, core-replaced, rejected by the saturated-star gate, and clip-corrected.

6 Background and diffuse emission

Two different quantities. The catalogues carry up to two background columns and they do not measure the same thing. `local_bkg` is the sigma-clipped annulus median that was handed to the PSF fitter, evaluated on whatever image that stage fitted; in the later stages that image has already had a smoothed background removed, so the column measures the residual of that subtraction and tends towards zero. `mean_modelsub_bkg` is the mean over a 3×3 pixel box on the per-exposure residual after only the star model is subtracted, combined across exposures as a sigma-clipped weighted mean. Writing D_i for the exposure data, S_i for the smoothed background subtracted before fitting, B_i for the annulus background and P_i for the fitted PSF model,

$$A_i = D_i - S_i - B_i, \quad R_i = D_i - S_i - P_i, \quad (1)$$

where A_i is the array the fitter minimised against and R_i is the image this column is measured on. Because R_i removes only the star, the second quantity is invariant across pipeline stages while the first is not — which is why they are plotted separately and never averaged together.

6.1 Distributions

Figure 6 shows both estimators. The median annulus background is 0.018 image units in F322W2. The residual-footprint background is absent from every catalogue here: it was introduced in July

2026 and these products predate it. Re-running the cataloguing stage would add it, and would make the background measurement independent of which pipeline stage produced the catalogue — which the annulus estimator is not. Towards the bright end the annulus estimator rises: the brightest one per cent of sources sit on a background that is 0.516 image units above the field median in F322W2, the largest such offset in this field. This is the star’s own wings entering its own background aperture, and it is a systematic rather than noise — it biases in one direction and grows with brightness. Any use of the background column for bright stars should account for it.

m4: background estimators




Figure 6: Background estimators for m4. First panel: the distribution of the annulus background handed to the fitter (solid) and, where available, the residual-footprint background (dashed). Remaining panels: the running median of each against brightness, with the 16–84th percentile band for the annulus estimator; the vertical dotted line marks the brightest 1 per cent, where a star’s own wings begin to enter its own background aperture. The residual-footprint column is present for 0 of 1 filters; it was added in 2026-07, so catalogues produced before then carry only the annulus estimator.

6.2 Spatial structure

Figure 7 is the test of the physical claim. The left column is a map of the diffuse emission built entirely out of the star catalogue; the right column sets the same per-source value against the drizzled mosaic sampled at each star. The rank correlation between them is 0.535 in F322W2. That is a real but partial correlation. The per-source background is picking up the extended emission, but a substantial part of its variance comes from something else — most plausibly the fitting residuals of neighbouring stars, which in a crowded field contribute to the same aperture.

7 Implications

- Positions in this field are repeatable to $3.2''$ per exposure at the bright end, so a proper-motion programme against an earlier epoch is limited by the epoch baseline and by the reference frame, not by these measurements.

m4: background structure vs diffuse emission




Figure 7: Spatial background structure in m4. Left column: the median per-source background in 56×56 sky cells – a map of the diffuse emission built entirely from the star catalogue. Right column: the same per-source value against the surface brightness of the drizzled mosaic sampled at that source’s position, with the running median in red and the Spearman rank correlation quoted. A high rank correlation is the evidence that the per-source background is measuring the astrophysical extended emission rather than a fitting artefact.

- The per-source background column is usable as a diffuse-emission tracer in its own right, sampled wherever there is a star; that is a spatial sampling no mosaic-based measurement provides, because it is unaffected by the stellar light that dominates the mosaic in these fields.
- Local registration residuals reach $335 \text{ m}''$ at the 95th percentile of tiles — widespread rather than isolated. Resolve this before using the affected bands for anything positional at the tens of milliarcsecond level.
- Re-cataloguing this field would add the stage-invariant residual-footprint background, which the current products predate; until then the background column here cannot be compared with a field catalogued after July 2026.

8 Reproducibility

Every figure in this document was produced by

```
python scripts/analysis/make_diagnostic_writeup.py --field m4
```

against `jwst_gc_pipeline` version `0.1.dev3+g737c2722d`. The numerical values quoted in the text are stored alongside the figures in `measurements.json`, so the prose and the plots are generated from one computation and cannot disagree.