From raw PS1 to the TESS template

Five-FFI stage trace and control-selection investigation

2 October 2026

Question. Do templates acquire an unexplained bias around removed bright stars, and at which processing stage does it appear?

Measured result. Fresh PS1-to-removal replay agrees exactly with the published centre cell in all five fields. Gaussian convolution, direct interior binning, template assembly and application of each saved Hotpants kernel reproduce their saved outputs at exact or floating-point precision.

Interpretation. The central depression in a star-minus-control stack is already present after source removal and is propagated by the later operators. Selecting unchanged empty sky creates a similar negative stack. More comparable controls remove much of the real-star contrast but leave smaller field-dependent residuals.

Conclusion: no additional unexplained downstream loss was found in the sampled interiors of the five original bootstrap templates. This does not establish unbiased photometry. The remaining population contrast needs a known-flux counterfactual test.

Coverage. Five FFIs; 45 spatially distributed skycells; 199,540,467 freshly replayed removal pixels; 3,124,925 interior OS4 output-pixel comparisons; 135 Gaussian cutout checks; 2,145 final-template patches, including 526 removed-star centres.

Boundaries. This audit covers the original bootstrap family, including C1's corrected WCS. It does not certify skycell-overlap seams, the separate v3 rebuild, the scientific adequacy of a fitted kernel, or the stellar flux model.

Frozen results: `/astro/armin/koji/syndiff/dev_runs/epsf_stage_trace_20261002/multifield/`.
The PDF contains embedded image plates and equations; no external images are needed to read it.

Contents

1 Dataset, inputs and coordinate conventions **2**

1.1 Sampling chosen before measuring the results . . . 2

1.2 Pixel origins and oversampling . . . 2

2 The processing chain and the direct tests **3**

2.1 Raw PS1 conversion and band combination . . . 3

2.2 Segmentation and bright-star removal . . . 3

2.3 Discrete Gaussian convolution . . . 3

2.4 Direct binning, masking and assembly . . . 4

2.5 Application of the saved Hotpants kernel . . . 4

2.6 How numerical residuals are summarized . . . 4

3 Direct numerical results **5**

4 Why a star-minus-control depression is not a photon-loss measurement **8**

4.1 The estimator used in the previous atlas . . . 8

4.2 A known counterfactual is needed to identify missing light . . . 8

4.3 Untouched empty-sky null . . . 9

4.4 Controls matched by the observed Gaia environment . . . 9

4.5 Spatial bootstrap and its scope . . . 9

5 Population results and their limits **10**

5.1 A common magnitude interval . . . 10

5.2 Complete field-by-magnitude values . . . 10

6 Corrections, confidence and conclusion **14**

6.1 What was corrected or checked . . . 14

6.2 The confidence is different for different claims . . . 14

6.3 Coverage boundaries . . . 14

6.4 The next discriminating experiment (proposed, not run) . . . 15

A Detailed image plates for each field **16**

B Reproducibility and records **57**

1 Dataset, inputs and coordinate conventions

The five selected full-frame images (FFIs) are listed in Table 1. Each audit uses the field's original bootstrap products. C1 uses the corrected-WCS FFI and rebuilt bootstrap; its actual input header was required to contain `WCSREFIT`. F2's previously registered glint rectangle, $x < 640$ and $y < 448$ native science pixels, is excluded from both the real-star and random-control populations. The numerical template checks themselves do not depend on a science-background fit.

Field	Sector/camera/CCD	FFI stem	Important note
F1	24/2/2	`tess2020120182919-s0024-2-2`	Original bootstrap
F2	20/3/3	`tess2020013215923-s0020-3-3`	Glint exclusion
S22	22/3/3	`tess2020071012921-s0022-3-3`	Original bootstrap
C1	24/1/2	`tess2020118235919-s0024-1-2`	Corrected WCS
C4	23/4/2	`tess2020087042920-s0023-4-2`	Common-protocol reference

Table 1: The selected FFIs. These names identify the camera and CCD as well as the time stem.

1.1 Sampling chosen before measuring the results

Nine cells per field were selected as the master-map owners of the 3×3 detector grid at $x, y \in \{384, 1024, 1664\}$ native science pixels. The original PS1 r, i, z, y data, masks and uncertainty maps were freshly retrieved for the centre cell of each field. Source removal was replayed on its complete array. The other eight sampled cells use recipe-matched published combined inputs. No failed sample was replaced on the basis of its measured value.

The combined recipe is `e17a198a4942aa2d`, with `footprint_v1` bright-star removal and a Gaia projection catalogue limited to RP magnitude 18. The frozen processing revision is `648c0bc`. The recorded band weights are $0.254, 0.4368, 0.1654, 0.1438$ for r, i, z, y respectively; they are this dataset's recipe, not a substitution of current site defaults.

1.2 Pixel origins and oversampling

The native science image is the full-FFI slice $[0 : 2048, 44 : 2092]$. Thus

$$x_{\mathrm{sci}} = x_{\mathrm{FFI}} - 44, \qquad y_{\mathrm{sci}} = y_{\mathrm{FFI}}, \qquad j_x = 4(x_{\mathrm{FFI}} - 36 + \tfrac{1}{2}) - \tfrac{1}{2}, \quad j_y = 4(y_{\mathrm{FFI}} + 8 + \tfrac{1}{2}) - \tfrac{1}{2}. \tag{1}$$

$x_{\mathrm{FFI}}, y_{\mathrm{FFI}}$ are zero-based full-FFI pixel-centre coordinates returned by the input WCS; $x_{\mathrm{sci}}, y_{\mathrm{sci}}$ are zero-based native science coordinates. j_x, j_y are OS4 pixel-centre indices. The factor 4 is the frozen oversampling factor; the template begins at full-FFI bounds $(36, -8)$, eight native pixels outside the science bounds. These bounds come from the template header. Catalogue positions use the dataset-private FFI WCS and epoch-2016 catalogue coordinates, consistently with the source-removal bookkeeping. The earlier C4-only atlas-refit/proper-motion positions are not mixed into the common five-field comparison.

Each field's FFI, template, `hp_c` and kernel files were SHA256-hashed before and after its audit. The hashes were unchanged. Every contributing cell has a recorded combined fingerprint and removal-record hash.

2 The processing chain and the direct tests

2.1 Raw PS1 conversion and band combination

On each available original band pixel the pinned conversion is

$$L_\beta(p) = \frac{o_\beta + 2s_\beta \sinh[R_\beta(p)/a]}{\tau_\beta}, \qquad I_0(p) = \sum_{\beta\in\{r,i,z,y\}} w_\beta L_\beta(p), \qquad a = \frac{2.5}{\ln 10}. \tag{2}$$

p indexes a native PS1 pixel; β labels a band. R_β is its stored logarithmic PS1 value. o_β, s_β and τ_β are the original FITS-header values `BOFFSET`, `BSOFTEN` and `EXPTIME`; the first two are in the PS1 linear-flux convention and the last is in seconds. a is the dimensionless conversion constant in the pinned code. L_β is the resulting linear PS1 flux rate per pixel. w_β is the dimensionless recipe weight listed above. I_0 is the combined PS1 image in its own linear flux-rate units per PS1 pixel; those units must not be assumed to equal TESS electron/s. Missing-data patterns, input masks and uncertainty conversion are propagated by the pinned routine. Unavailable bright-star cores in the image plates are shown as unavailable, not interpreted as astrophysical zero sky.

2.2 Segmentation and bright-star removal

The algorithm first sets non-segment, non-bright pixels to zero. It constructs 8-connected footprints from detected segments, the bright-core mask and the documented saturated-pixel flags. Gaia sources brighter than catalogue TESS magnitude 13 trigger footprint removal, with the recorded removal radius cap. A footprint may also contain faint neighbours, which are recorded where catalogued.

$$I_{\rm seg} = S(I_0), \qquad I_1 = Z(I_0). \tag{3}$$

S is the pinned non-segment/background-zeroing operation; $I_{\rm seg}$ is its intermediate image. Z is the complete segmentation and footprint-removal operation with the frozen Gaia catalogue, masks and thresholds; I_1 is the published combined image after removal. Both images retain the units of I_0. Z is replayed with the recorded inputs, not fitted to reproduce a result.

The full-array raw-removal replay is a provenance and algorithm-reproduction test. It cannot determine the unobserved faint-source light under a masked bright-star core. Pixels intentionally suppressed by Z are a scientific choice that downstream modelling must account for; agreement of the replay does not by itself certify that choice as bias-free.

2.3 Discrete Gaussian convolution

An independent direct calculation uses the actual normalized discrete kernel:

$$g(u) = \frac{\exp[-u^2/(2\sigma^2)]}{\sum_{v=-r_G}^{r_G} \exp[-v^2/(2\sigma^2)]}, \qquad G(u,v) = g(u)g(v), \qquad C(p) = \sum_q G(q) I_1(p-q). \tag{4}$$

u, v are integer PS1 pixel offsets in $[-r_G, r_G]$; $\sigma = 40$ PS1 pixels and $r_G = 470$ PS1 pixels are the frozen Gaussian width and truncation radius. q denotes a two-dimensional offset in that support. G is dimensionless and unit-sum; C is the canonical convolved cell, in the same per-PS1-pixel units as I_1.

Three bright-star cutouts per sampled cell were checked. Each measured 361×361 window lies beyond the adjacent-cell 480-pixel overlap plus the Gaussian support. This ensures the direct calculation and canonical output refer to the same input pixels. An initial C4 window that did not satisfy that same-input condition was corrected by moving the test domain, not by relaxing a numerical tolerance.

2.4 Direct binning, masking and assembly

$$\widehat{T}_j = \sum_{p:\, a(p)=j} [1 - M(p)]\, \widetilde{C}_{\mathrm{pad}}(p), \qquad E_{B,j} = \widehat{T}_j - T_{\mathrm{contrib},j}, \qquad E_{A,j} = T_{\mathrm{contrib},j} - T_j. \tag{5}$$

j indexes an OS4 output pixel; $a(p)$ is the frozen PS1-to-OS4 assignment. C_{pad} is the saved padding-aware convolved input. The tilde denotes the pipeline's zero treatment of unavailable inputs during summation, not a physical estimate of the sky. $M(p)$ is 1 when mask bit 12 is set and 0 otherwise. $\widehat{T}_j$ is the independently binned prediction; $T_{\mathrm{contrib},j}$ is the saved cell contribution; T_j is the published assembled template. All three sums and residuals are in PS1-combination flux units per OS4 pixel.

The comparison domain has positive pixel counts equal between the frozen map and the saved contribution, eroded by eight OS4 pixels around count/support changes. These are pure-cell interior checks. The spatially varying assignments at seams and the reconstruction of padding itself are not certified by this test. The actual bit-12 loss is included explicitly in the prediction; the comparison does not silently assume all convolved pixels contribute.

2.5 Application of the saved Hotpants kernel

$$\widehat{H}_n = \sum_{j \in \mathcal{B}(n)} [K_\theta T]_j + b_n, \qquad E_{H,n} = \widehat{H}_n - H_n. \tag{6}$$

n indexes a native output pixel; $\mathcal{B}(n)$ is its aligned 4×4 OS4 block. K_θ is the spatial convolution with the field's saved kernel coefficients θ, including the fitted flux-unit conversion. b_n is the saved background polynomial, constant for this recipe; H_n is the published `hp_c` image. $\widehat{H}_n$ and $E_{H,n}$ are in electron/s per native pixel. Kernel basis, spatial polynomial coordinates, block boundaries and background indexing match the saved fit. No coefficients, background levels or flux scales are refitted for this replay.

The allocated coefficient array contains unused tail entries. Reading its final allocated element as background would be incorrect; the populated background coefficient for the recorded layout is at index 722. The replay uses that verified layout. A saved-kernel application test checks application of a fit, not whether that fit is the ideal PSF match to the astrophysical sky.

2.6 How numerical residuals are summarized

$$R_{\mathrm{rms}} = \left(\frac{1}{N_P} \sum_{p \in \mathcal{P}} E_p^2 \right)^{1/2}, \qquad R_{\max} = \max_{p \in \mathcal{P}} |E_p|. \tag{7}$$

$\mathcal{P}$ is the stated valid comparison set; N_P is its pixel count; E_p is the relevant stage's prediction-minus-published residual, in that stage's units. R_{rms} and $R_{\max}$ are respectively its root-mean-square and maximum absolute residual. Pixel counts describe tested coverage, not a number of statistically independent astrophysical measurements.

3 Direct numerical results

The five complete centre-cell source-removal replays reproduce 199,540,467 output pixels exactly. The later tests cover 45 distributed skycells, 135 strict-interior Gaussian windows, 3,124,925 interior OS4 pixels and 2,145 `hp_c` patches, including 526 removed-star centres. There is no fitted residual minimization in these checks. The worst

Field	Interior OS4 pixels	Binning RMS PS1 units/OS4 pixel	`hp_c` replay RMS electron/s/native pixel
F1	625,191	2.03×10^{-6}	7.97×10^{-7}
F2	642,696	2.02×10^{-6}	8.01×10^{-7}
S22	593,454	1.92×10^{-6}	7.69×10^{-7}
C1	622,684	1.77×10^{-6}	7.23×10^{-7}
C4	640,900	2.30×10^{-6}	9.37×10^{-7}

Table 2: Direct operator comparisons. Saved contributions equal the assembled template exactly on these same interior pixels.

Gaussian-window RMS is 1.66×10^{-8} PS1 units per PS1 pixel. The maximum final `hp_c` replay difference is 3.05×10^{-5} electron/s per native pixel. The residuals in the image plates have no coherent loss pattern on their displayed narrow scales. Their sizes are consistent with floating-point calculation and narrowing, rather than a new astrophysical component.

What follows. In the tested interiors, the later operations reproduce their actual inputs and saved outputs. The depression is already present in the source-removal output; Gaussian smoothing broadens it.

What does not follow. These equalities do not imply that zeroing a footprint preserves all neighbour light, that the saved kernel is scientifically exact, or that a fitted star flux is unbiased. Those questions require a known sky or a controlled injected-source counterfactual.

Raw PS1 evidence across the five FFIs

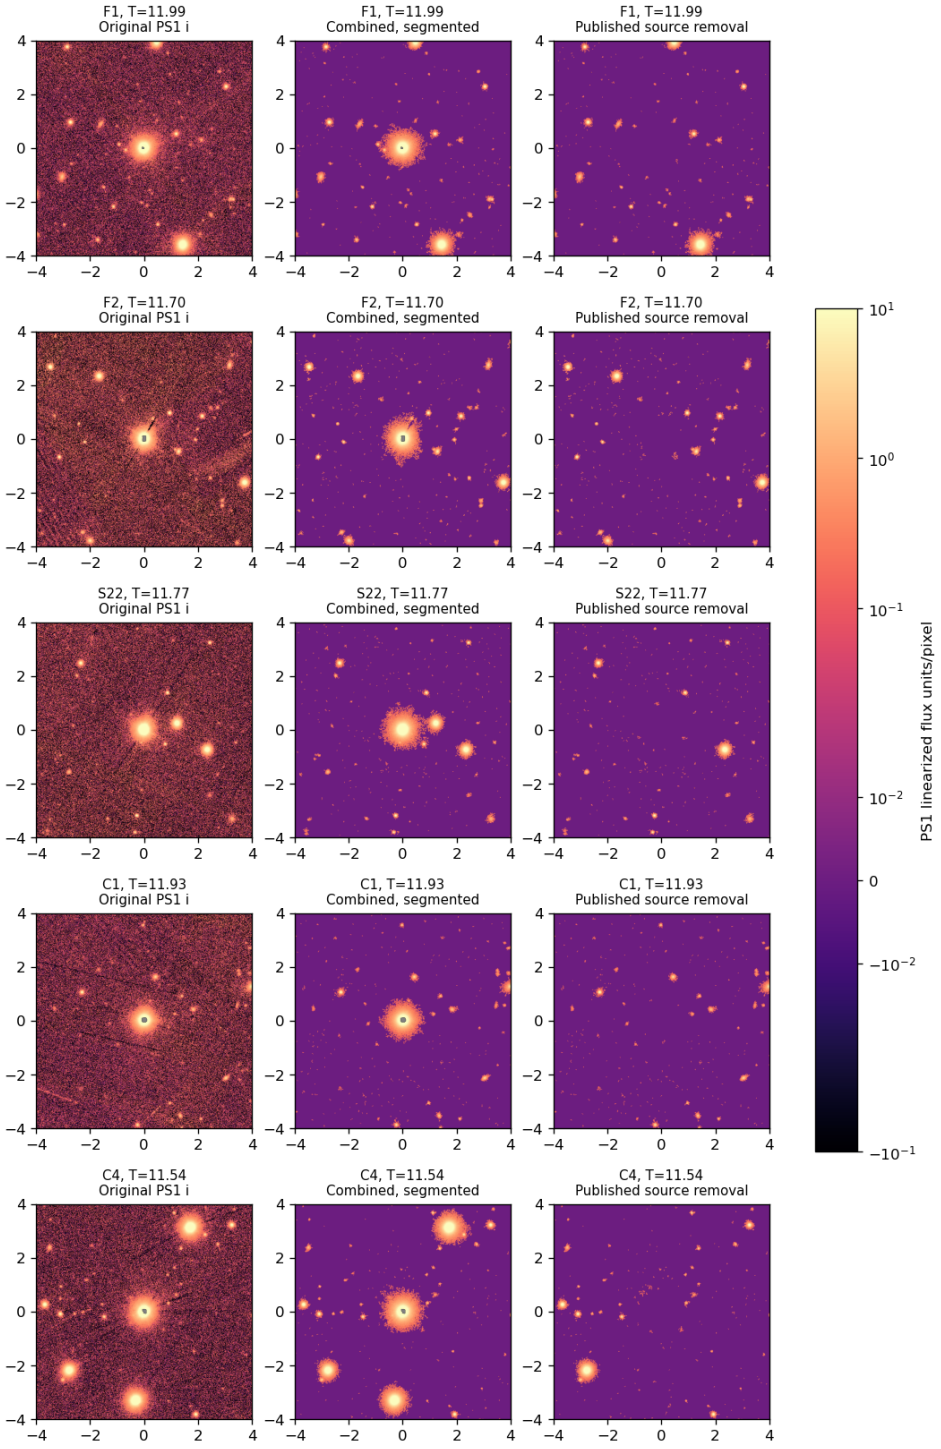


Figure 1. The same source positions before and after removal. Left: original PS1 i band; middle: segmented band combination; right: published source-removal output. Grey pixels are unavailable original PS1 data. The full $r/i/z/y$ examples appear in the field appendix. These images show what is suppressed; the raw-replay equality was checked on complete centre-cell arrays.

Actual final-template pixels and their replay

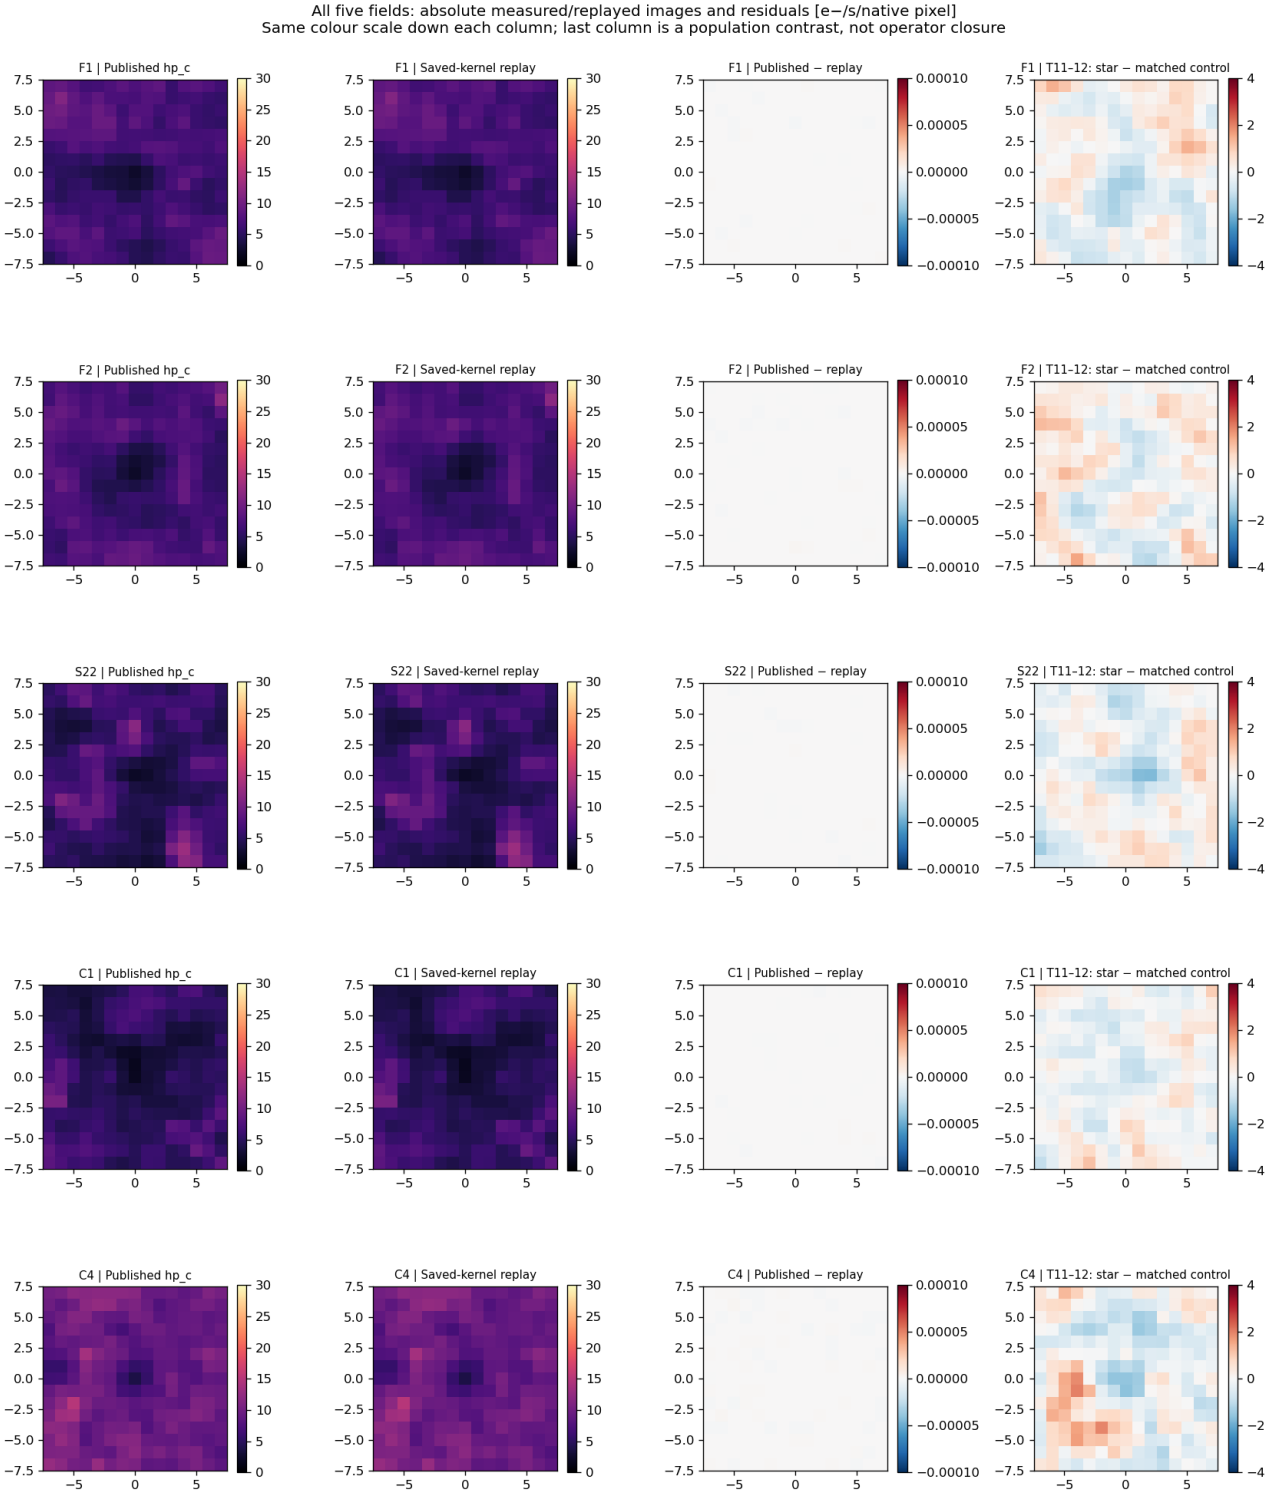


Figure 2. Each row is one field. The first two columns are absolute published and saved-kernel-replayed mean images at removed-star positions. The third is their difference, on a narrow common scale. The last is a separate TESS-magnitude 11–12 population contrast against matched controls. The operator residual is numerically negligible while the population contrast may remain; the two columns measure different things.

4 Why a star-minus-control depression is not a photon-loss measurement

4.1 The estimator used in the previous atlas

The earlier statistic removes each stamp's local annulus level, sums a core aperture and compares star locations with controls. It is reproduced here without own-flux normalization or the former relative-isolation cut.

$$\mu(H, \boldsymbol{x}) = \frac{1}{|\mathcal{A}(\boldsymbol{x})|} \sum_{n \in \mathcal{A}(\boldsymbol{x})} H_n, \tag{8}$$

$$A(H, \boldsymbol{x}) = \sum_{n \in \mathcal{Q}(\boldsymbol{x})} [H_n - \mu(H, \boldsymbol{x})], \tag{9}$$

$$D = \frac{1}{N_S} \sum_{s=1}^{N_S} \left[A(H, \boldsymbol{x}_s) - A(H, \boldsymbol{x}_{c,s}) \right]. \tag{10}$$

$\boldsymbol{x}$ is a crop-local native-pixel position; $\mathcal{A}(\boldsymbol{x})$ is the 20–30-native-pixel annulus around its subpixel position. μ is an arithmetic mean in electron/s/pixel. $\mathcal{Q}$ is the set within radius 3 of the rounded stamp centre, matching the implemented historical aperture convention; its strict integer-centred radius cut includes 25 native pixel centres. No additional continuous πr^2 area factor is applied to the reported sums. A and D have units electron/s. N_S counts selected stars or selected random positions; $\boldsymbol{x}_s$ is the selected position and $\boldsymbol{x}_{c,s}$ its control. Original controls are 25 native pixels east. The angle-bracket notation used in earlier reports means this same arithmetic ensemble mean.

4.2 A known counterfactual is needed to identify missing light

Let an ideal template contain exactly the retained astrophysical light that the scientific model intends to subtract. Then

$$d_s = A(H, \boldsymbol{x}_s) - A(H, \boldsymbol{x}_{c,s}) = \epsilon_s + \left[A(H_{\text{ideal}}, \boldsymbol{x}_s) - A(H, \boldsymbol{x}_{c,s}) \right], \qquad \epsilon_s = A(H - H_{\text{ideal}}, \boldsymbol{x}_s). \tag{11}$$

H_{ideal} is the counterfactual ideal template in electron/s/native pixel; it is not supplied by this audit. d_s is one observed star-control aperture contrast. ϵ_s is the template-model error in the star aperture, in electron/s. The term in square brackets is the differing source environment between the star and control apertures. It is generally nonzero, including for stars with zero recorded deleted neighbours. Both terms contribute to D.

Thus a negative D, a round dark stack or a weak correlation of D with modelled deleted-neighbour flux does not uniquely identify ϵ_s. The previous per-star "closure" used an unmatched population statistic, rather than a direct before/after photon budget. The precise ideal sky beneath an unavailable PS1 bright-star core cannot be reconstructed by replacing NaNs with zeros.

4.3 Untouched empty-sky null

Random test positions and their east-shifted controls are both more than 5 native pixels from a bright Gaia source. No image is changed. The selected position is required to be at least a specified distance from every retained Gaia source with catalogue TESS magnitude between 13 and 18; the nearby control is unrestricted by that retained-source distance. Each setting uses 4,000 positions. This tests whether changing the source population alone can create a central negative stack.

Field	No exclusion	No retained source within 2 pixels
F1	0.8 ± 4.6	-54.8 ± 4.2
F2	-22.6 ± 27.3	-127.9 ± 62.2
S22	-2.7 ± 4.1	-46.2 ± 3.4
C1	-2.5 ± 3.9	-44.2 ± 3.3
C4	-0.9 ± 5.6	-85.1 ± 4.7

Table 3: Untouched-sky aperture contrasts in electron/s. Errors are spatial-block bootstrap uncertainties. These are null demonstrations, not corrections fitted to real stars.

The negative central feature appears without a bright-star removal at the chosen centre. The F2 control stamps contain conspicuous off-centre features and correspondingly large bootstrap errors. Its scalar curve should not be treated as a clean universal amplitude; the images are part of the evidence. The null disproves the uniqueness of the earlier missing-pedestal interpretation, but does not prove that all real-star contrasts are selection.

4.4 Controls matched by the observed Gaia environment

For each real star, four random control positions are selected in the same 128×128 native-pixel detector cell and the same 0.2-native-pixel bin of distance to the nearest retained Gaia source. Both real and control locations are away from other bright stars; original east-shifted controls are required to satisfy that bright-source condition too. Matching uses catalogue positions, never the central template intensity. The displayed real-star subset has zero recorded deleted neighbours.

$$D_{\mathrm{match}} = \frac{1}{N_S} \sum_{s=1}^{N_S} \left[A(H, \boldsymbol{x}_s) - \frac{1}{4} \sum_{\ell=1}^{4} A(H, \boldsymbol{x}_{c,s,\ell}) \right]. \tag{12}$$

D_{match} is the matched mean aperture contrast in electron/s; ℓ indexes the four matched controls per star. All other symbols and the aperture convention are defined in Eq. 10. Retained-source brightness, the full neighbour configuration and catalogue completeness are not matched. Consequently this experiment reduces an important confound without creating an exact ideal-sky counterfactual.

4.5 Spatial bootstrap and its scope

For each resampled detector cell the star contributions are kept together:

$$D^{*(t)} = \frac{\sum_c w_c^{(t)} \sum_{s \in c} d_s}{\sum_c w_c^{(t)} N_c}, \qquad \widehat{\sigma}_D = \mathrm{sd}_{t=1,\ldots,500}\{D^{*(t)}\}. \tag{13}$$

c indexes occupied 128-pixel detector cells; N_c is their selected-star count. d_s is the relevant unmatched or matched per-star contrast. $w_c^{(t)}$ is the number of times cell c is drawn with replacement in bootstrap realization t, and $D^{*(t)}$ is that realization's weighted mean, in electron/s. $\widehat{\sigma}_D$ is the standard deviation of 500 realizations. This preserves local spatial dependence better than a pixel-noise-only error. It does not include the systematic uncertainty of incomplete scene matching.

5 Population results and their limits

5.1 A common magnitude interval

For catalogue TESS magnitude 11–12, matching reduces the aperture contrast in every field: F2 is consistent with

Field	Stars	Original control	Matched control
F1	1,292	-42.6 ± 6.7	-16.9 ± 4.7
F2	1,256	-22.2 ± 7.0	-4.3 ± 7.8
S22	1,243	-32.2 ± 7.3	-16.7 ± 5.1
C1	1,128	-27.1 ± 6.1	-10.5 ± 4.2
C4	1,240	-42.6 ± 7.4	-17.6 ± 5.5

Table 4: Aperture contrasts in electron/s on the same retained star subset in each row. The image pairs and their differences are shown in Figure 3.

zero in this interval. F1, S22 and C4 remain negative at about 3.2–3.6 statistical standard errors; C1 is about 2.5. These are not systematic-error-inclusive or multiple-comparison-adjusted significances. Several residual stamps are broad, asymmetric or off-centre rather than a pure round central mode.

Interpretation: the residuals are worth following up, but their aperture sums cannot yet be equated with a stellar photometric error. A universal added template pedestal, a brightness-width term or a claimed fractional stellar-flux bias is not justified by this comparison alone.

5.2 Complete field-by-magnitude values

Field	TESS magnitude	Stars	Original contrast	Matched contrast
F1	8–10	249	-64.8 ± 15.7	-4.8 ± 10.4
F1	10–11	535	-35.0 ± 11.2	-11.7 ± 6.8
F1	11–12	1,292	-42.6 ± 6.7	-16.9 ± 4.7
F1	12–13	2,956	-22.2 ± 4.4	-9.2 ± 3.3
F2	8–10	264	-45.7 ± 11.6	-23.2 ± 8.8
F2	10–11	504	-90.8 ± 40.4	-32.2 ± 16.2
F2	11–12	1,256	-22.2 ± 7.0	-4.3 ± 7.8
F2	12–13	2,948	-19.9 ± 4.7	-2.4 ± 6.4
S22	8–10	300	-70.8 ± 11.9	-10.9 ± 6.6
S22	10–11	528	-34.0 ± 9.0	-15.4 ± 6.0
S22	11–12	1,243	-32.2 ± 7.3	-16.7 ± 5.1
S22	12–13	2,789	-22.7 ± 4.5	-5.2 ± 2.8
C1	8–10	275	-44.9 ± 11.0	-11.2 ± 5.9
C1	10–11	445	-31.5 ± 8.1	-8.4 ± 5.5
C1	11–12	1,128	-27.1 ± 6.1	-10.5 ± 4.2
C1	12–13	2,625	-15.2 ± 4.3	-5.5 ± 2.8
C4	8–10	208	-89.2 ± 18.4	-19.5 ± 10.8
C4	10–11	467	-61.8 ± 13.0	-34.3 ± 8.1
C4	11–12	1,240	-42.6 ± 7.4	-17.6 ± 5.5
C4	12–13	3,039	-37.8 ± 5.0	-15.5 ± 4.0

All contrasts are electron/s and all errors use the stated spatial bootstrap. The full absolute star/control stacks and both differences for every magnitude bin are provided in the field appendix, on shared colour scales.

The images behind the common-magnitude comparison

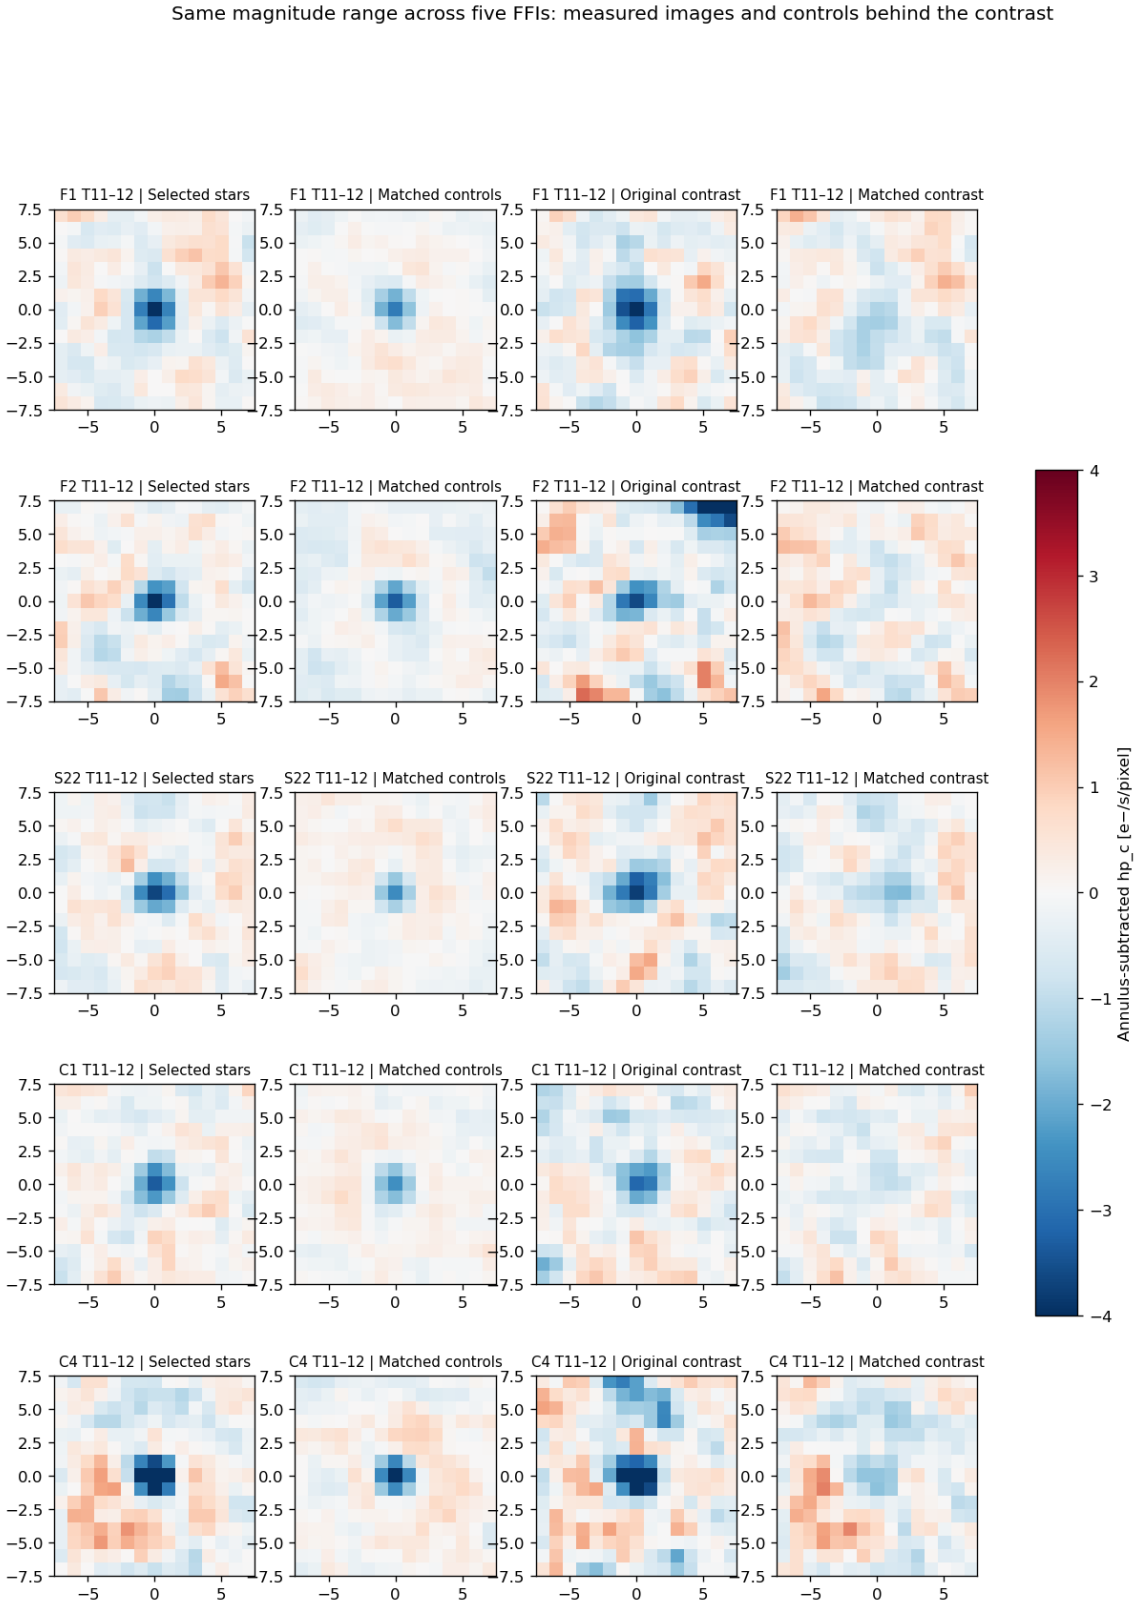


Figure 3. TESS magnitude 11–12 in each field. Columns: selected stars, matched controls, original star-minus-east-control contrast, and star-minus-matched-control contrast. The matched controls themselves have central depressions. The remaining contrasts are smaller, but some retain broad or lopsided structure. All panels share the same intensity scale.

Remaining contrasts across all fields and magnitude bins

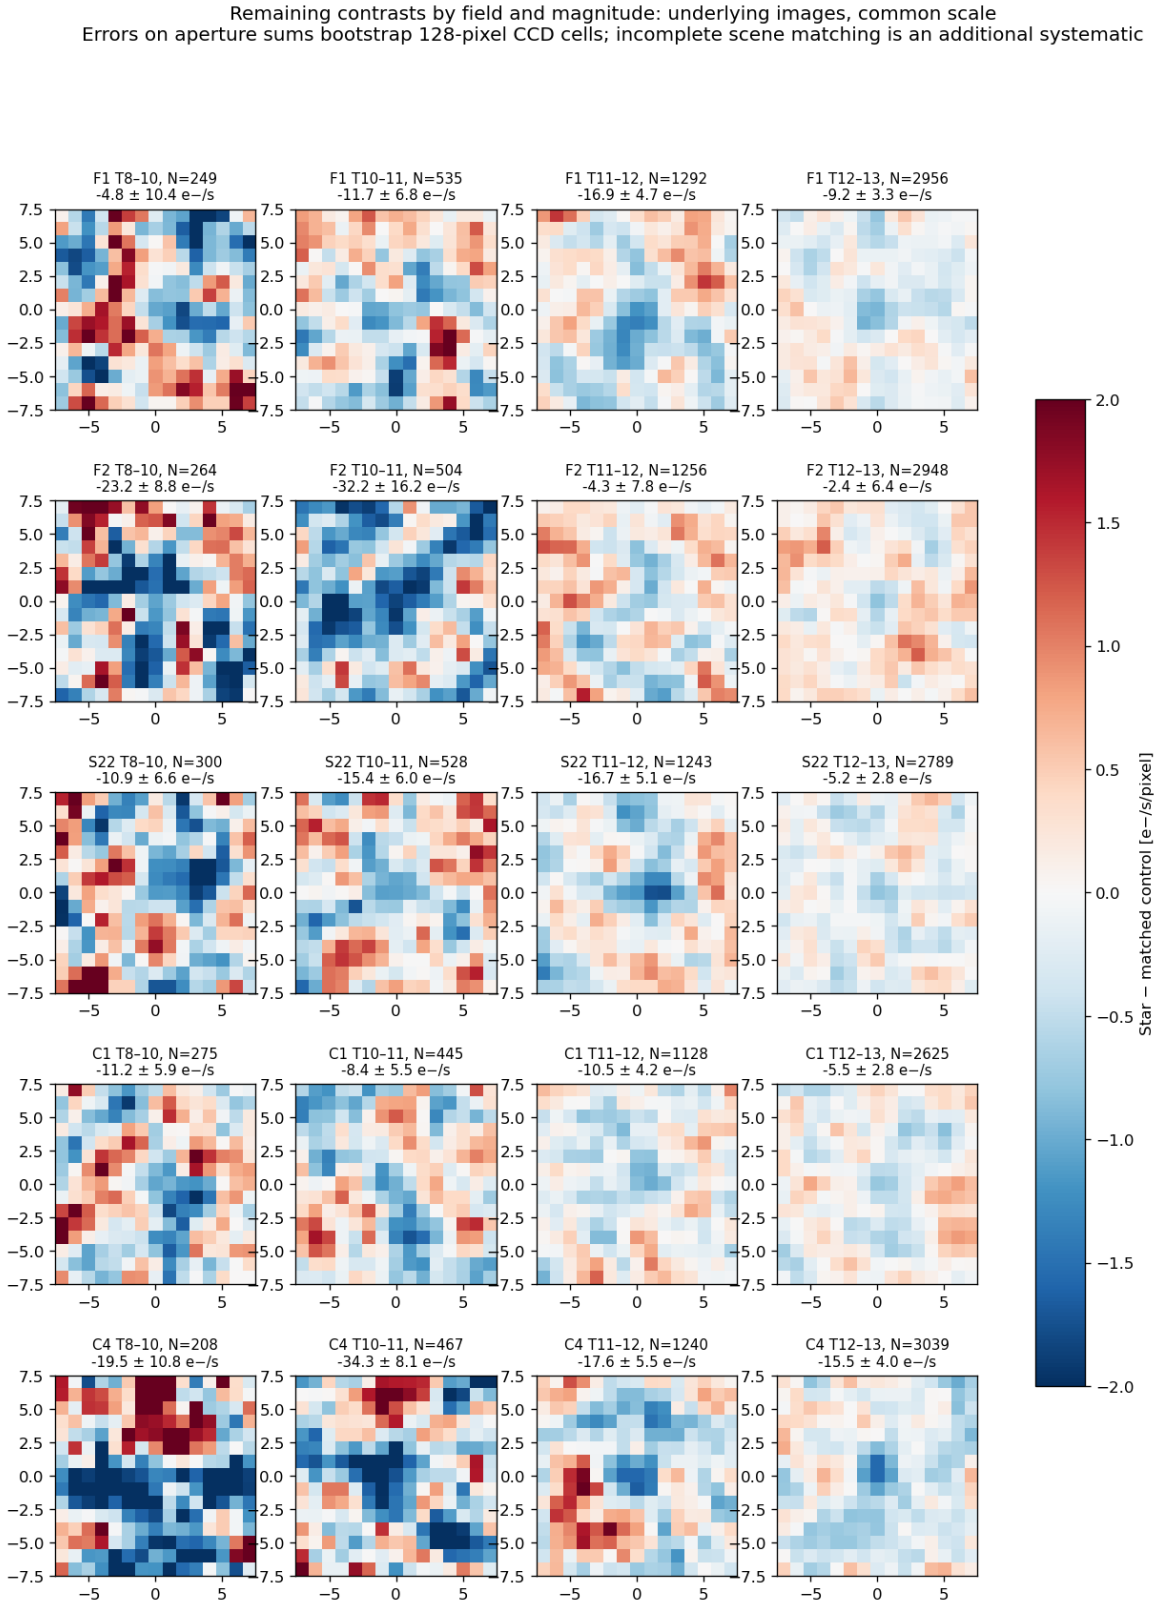


Figure 4. The measured difference images behind the complete numerical table. Every panel has the same colour scale; titles give the star count and measured aperture contrast. A small aperture sum is not by itself a decomposition into PSF width, background, missing-source light or flux error.

Summary plots read alongside the measured images

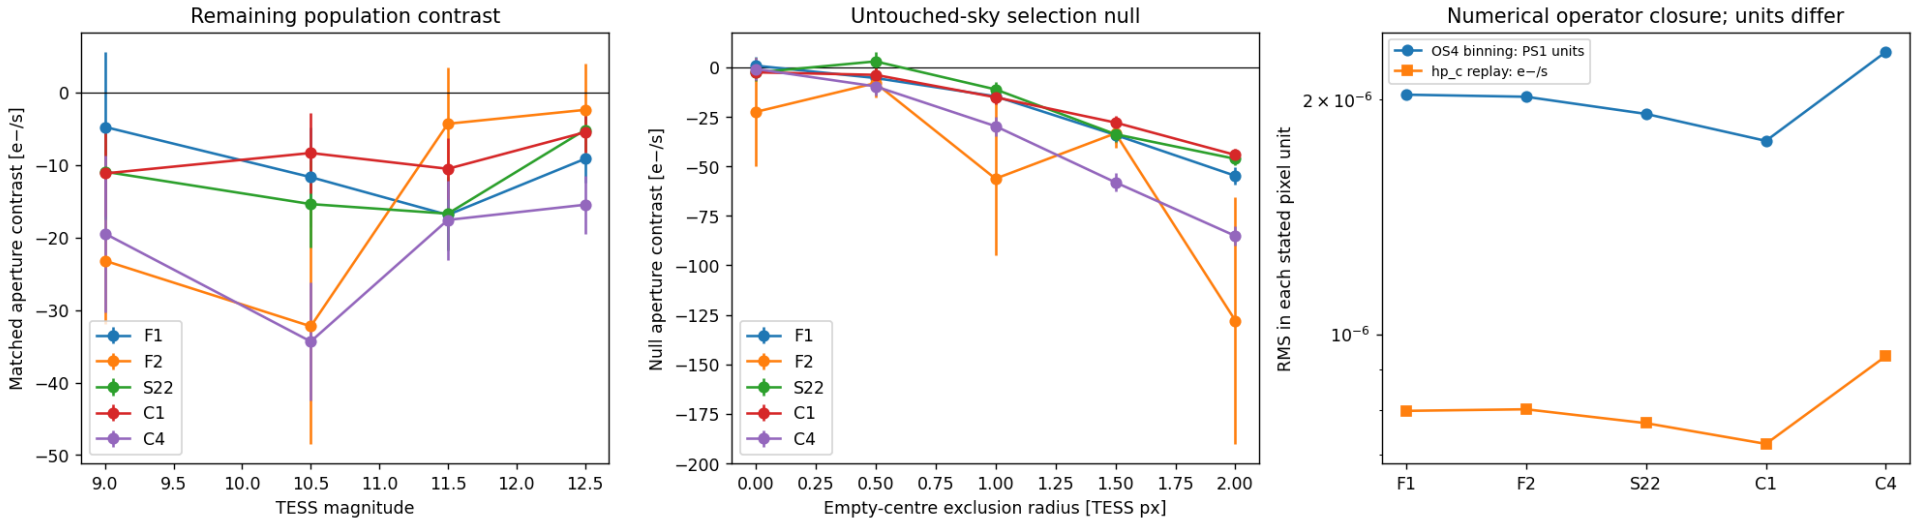


Figure 5. Left: field-by-magnitude matched aperture contrasts. Middle: the untouched-sky selection null. Right: numerical operator residuals, each in its stated unit; their units differ, so their vertical separation is not an astrophysical comparison. The raw/stage/control image plates are the underlying evidence for these curves.

6 Corrections, confidence and conclusion

6.1 What was corrected or checked

- The original relative-isolation cut conditioned neighbour light on target flux, which made faint stars sample emptier surroundings. This audit uses no such cut for its operator checks or control experiment.
- The earlier own-flux normalization, unpaired masks and different background-fit pixel sets could distort the shape of raw-versus-processed differences. The present stage checks use actual absolute pixel inputs and fixed operators rather than those fitted stack normalizations.
- The original Gaussian test admitted an adjacent-cell overlap through its support. The final same-input comparison excludes that overlap, the kernel support and the measured window. No numerical tolerance was weakened.
- Saved Hotpants coefficients have unused allocated tail slots. The background index was verified before accepting a replay result.
- Neighbour records are now associated across all cell/segment links before deduplicating source IDs, and neighbours without usable magnitudes still count as recorded neighbours. Four previously zero-labelled common stars in each of C4 and F2 have recorded neighbours; two per field involve missing magnitudes. This is too small to explain the broad stack result, but corrects the meaning of "zero recorded neighbours."
- C1's corrected WCS and the actual template bounds were checked. The five-field comparison uses one coordinate convention and each field's own saved Hotpants conversion, not a copied nominal C4 flux scale.

6.2 The confidence is different for different claims

High confidence within the tested domain: input provenance and reproduction of source removal, Gaussian convolution, interior binning/assembly and application of the saved Hotpants kernel. The comparisons are direct pixel calculations and reproduce all five fields.

Strong qualitative evidence: an unmatched star-minus-control depression is sensitive to source-population selection. The unchanged-sky null and the matched-control image pairs demonstrate this directly.

Unresolved: the size and physical cause of the smaller remaining population contrast, and whether a particular downstream stellar-flux estimator is biased. Incomplete scene matching and unavailable PS1 cores are not covered by the bootstrap errors.

6.3 Coverage boundaries

Only one complete centre skycell per field was freshly fetched and replayed from raw bands. The nine-cell later-stage audit is a spatial sample, not an exhaustive whole-detector certification. Eroded interiors exclude assignment seams; loading the saved padding-aware array is not an independent reconstruction of every cross-projection padding step. The saved-kernel replay does not independently validate the kernel's scientific fit to an ideal sky.

A separate v3 rebuild addresses vertical skycell-overlap ownership. At the time of this audit, its products were present for F1/S22/C4 under `paper_dataset_20261001/template_v3`, and for F2 under `vrow_validate_20261001/F2_bootstrap`; C1's v3 bootstrap was absent at inspection. These products were not mixed into the repeat of the original bootstrap audit. This document does not certify them or their seams.

No ePSF training, pipeline rerun, template modification or new fitted parameter was performed for this audit. The older ePSF-estimator validation failures are a separate investigation.

6.4 The next discriminating experiment (proposed, not run)

Inject a faint source of known flux into a matched raw PS1 scene and propagate paired with/without-injection differences through the same frozen operations. Record whether the injected source is removed with a bright-star footprint and whether the downstream scene model includes the recorded removed source. This supplies a known counterfactual photon budget, including the dependence on position, footprint and source brightness, that the present random-sky contrasts lack.

One useful definition, with the later operators and masks held fixed, would be

$$\Delta H(f) = K_\theta B[(1-M)G * \{Z(I_0 + fP) - Z(I_0)\}], \tag{14}$$

$$\Delta H_{\mathrm{ref}}(f) = K_\theta B[(1-M)G * (fP)], \tag{15}$$

$$\eta(f) = \frac{A(\Delta H(f), \boldsymbol{x}_{\mathrm{inj}})}{A(\Delta H_{\mathrm{ref}}(f), \boldsymbol{x}_{\mathrm{inj}})}. \tag{16}$$

f is a known injected integrated flux rate in the linear PS1 convention. P is an explicitly chosen PS1 source profile with unit sum on its pixel grid; hence fP has the same per-pixel units as I_0. $\boldsymbol{x}_{\mathrm{inj}}$ is its known native TESS position. ΔH is the processed with-minus-without-source change, in electron/s/native pixel. ΔH_{ref} is the change expected through the same later operators without source removal. Their constant saved background cancels. η is a dimensionless aperture-recovery ratio, defined only when the reference aperture response is nonzero. Z, G, B, M, K_θ and A retain the definitions and units given above.

This would measure the response to an actual known added source rather than the difference between two unrelated sky populations. The test should report whether the injected source is catalogued, retained, deleted with a bright footprint or absent from the downstream scene model. Source-removal recovery and stellar-flux-estimator bias are related but distinct: a known deleted source can be accounted for by a correctly specified scene model. No value of η has been measured in the present investigation.

Final conclusion. The C4 stage result generalizes to all five original bootstrap templates in the sampled interiors. No additional unexplained downstream loss was detected. Control selection explains a substantial part of the apparent depression, while smaller field-dependent contrasts remain. Actual stellar photometric bias is not established or excluded by this audit. The remaining inference requires the paired known-flux experiment, not a background pedestal inferred from the old contrast alone.

A Detailed image plates for each field

The appendix keeps absolute inputs, controls and differences together. The nine sampled-cell maps are split into groups of three cells so that their pixel structure and colour scales remain readable. Numeric closure is evaluated on the actual arrays; the plots are displays of those same pixels. Full-resolution images are embedded in the PDF and can be zoomed.

F1: all four original PS1 bands and source removal

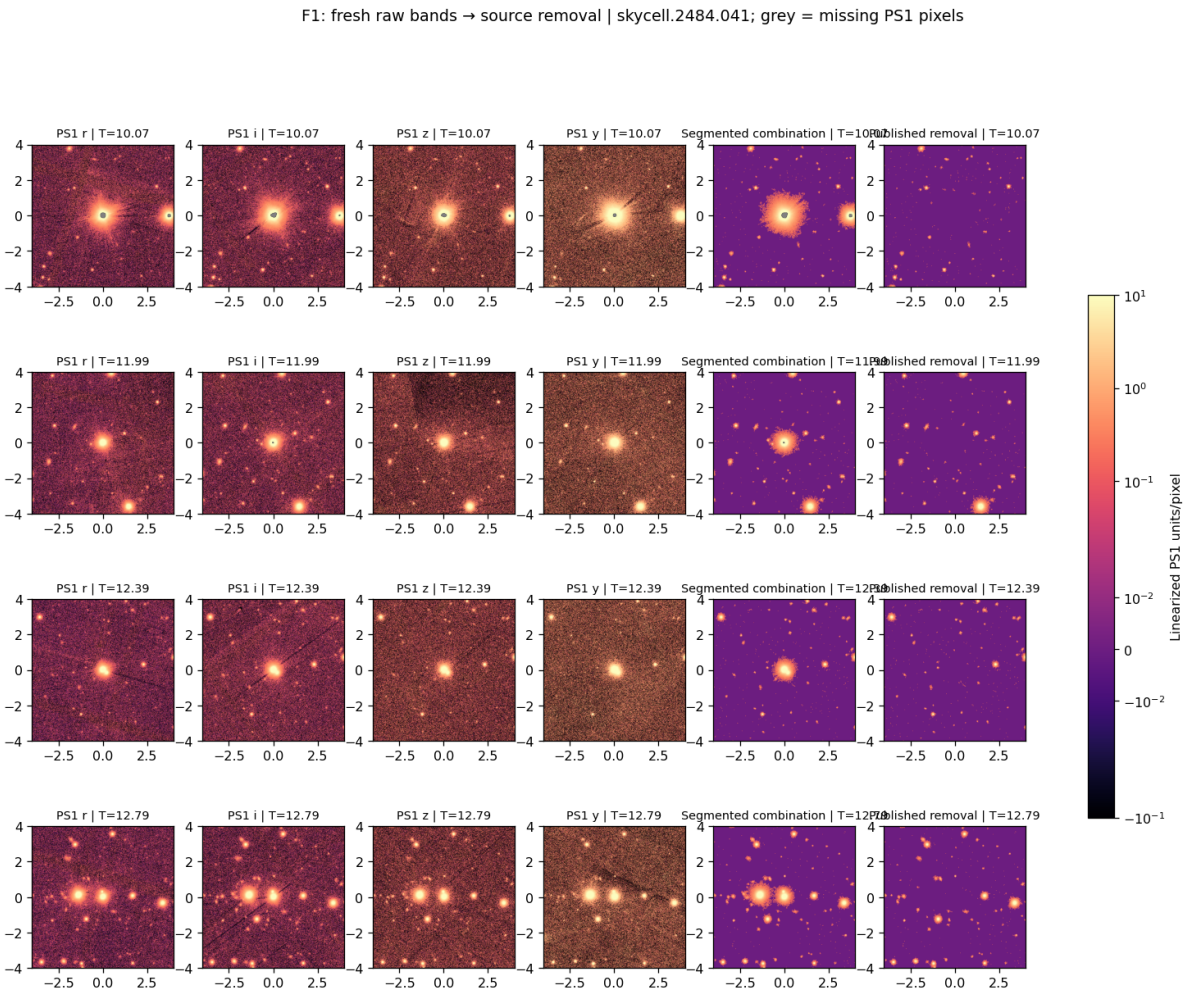


Figure 6. Four magnitude-spread examples, identical source positions. Columns show the original r, i, z, y bands, segmented combination and published removal. Grey marks unavailable original PS1 pixels; intensities in different bands differ physically. The centre-cell full-array replay is exact. The raw PS1 values have their own flux convention, not TESS electron/s.

F1: absolute stage images and paired differences

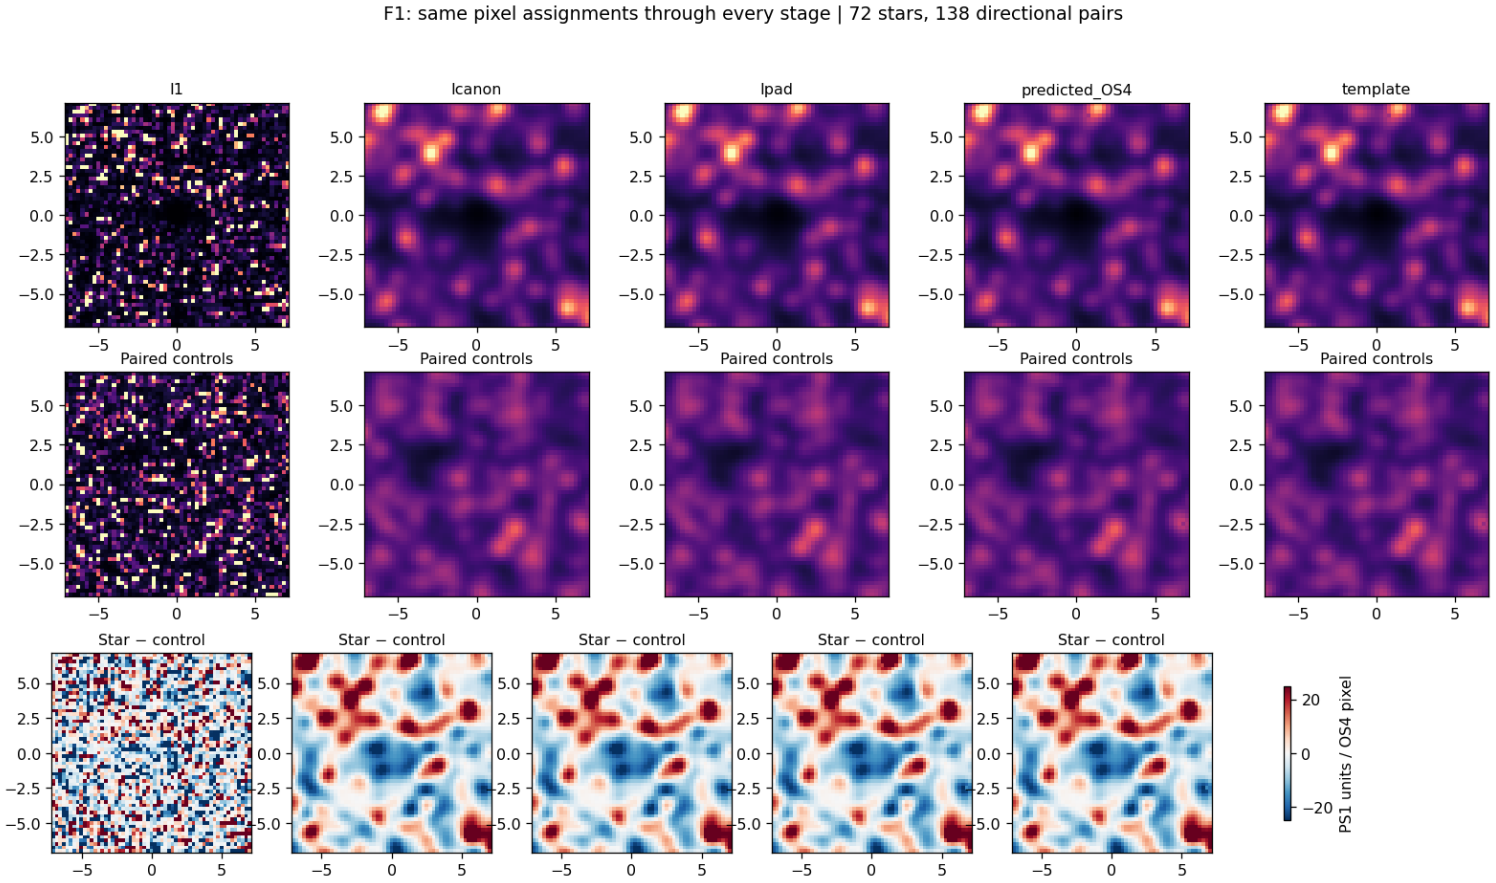


Figure 7. Top: mean at removed-star positions. Middle: paired directional controls. Bottom: their difference. Columns progress from published source-removed PS1 pixels through canonical and padding-aware convolution to predicted and published OS4 template. The central depression is already present in the first column; the later calculations propagate it. Native PS1 sums have been put on the same frozen OS4 assignments; units are PS1 flux per OS4 pixel.

F1: an individual patch and the mean saved-kernel replay

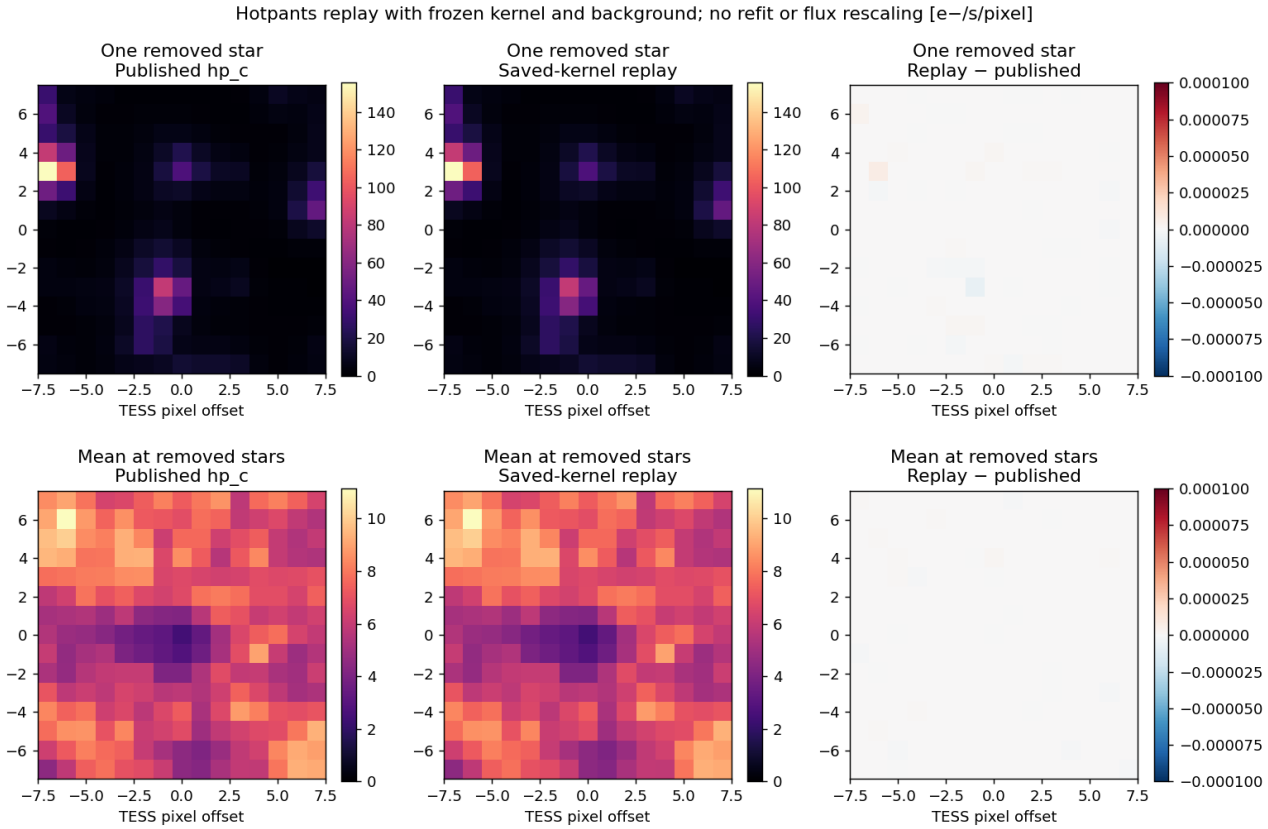


Figure 8. Published final template, replay with the frozen field kernel/background, and their difference. The first row is an individual removed-star patch; the second is the mean over tested star positions. The residual has its own narrow scale. No fit or flux rescaling was used; units are electron/s per native pixel.

F1: untouched-sky selection null

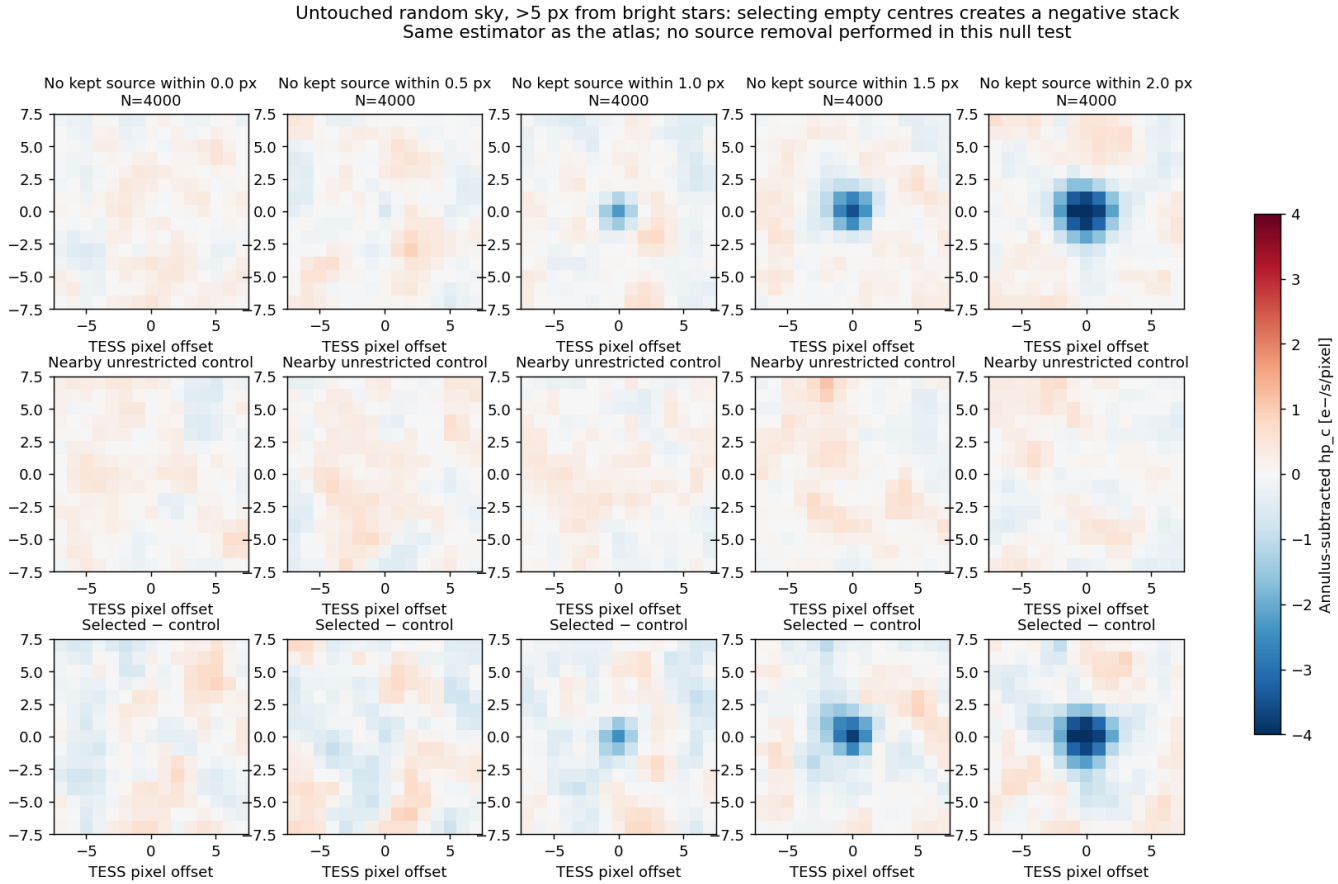


Figure 9. Columns change only the minimum distance to retained Gaia sources. Rows show selected unchanged locations, their nearby unrestricted controls and the difference. Both centres are away from bright stars. All images use the same scale in electron/s/native pixel after local annulus subtraction. A depression can be generated by changing the source population sampled by the estimator. Strong off-centre structures in the F2 controls account for its larger uncertainties.

F1: real stars, matched controls and both contrasts

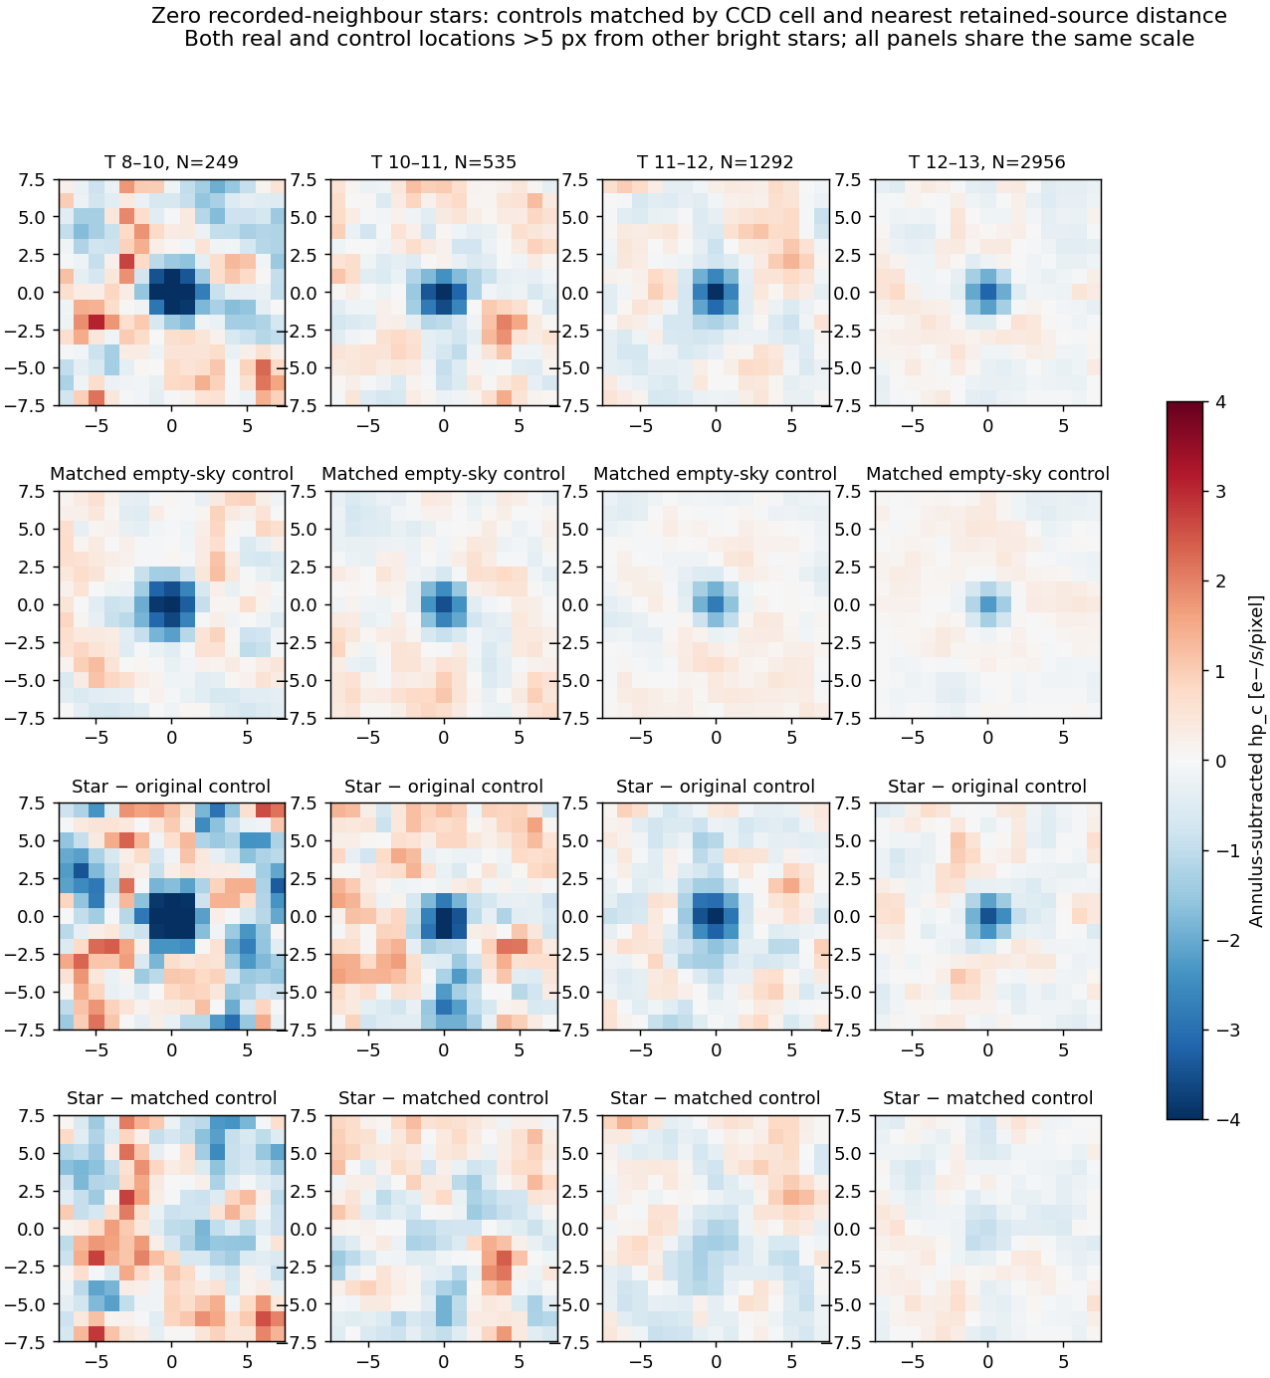


Figure 10. Columns are the four catalogue-magnitude bins. Rows show absolute annulus-subtracted star stamps, matched controls, the original contrast and the matched contrast. Matching never uses central template intensity. All panels share one scale, in electron/s/native pixel. The complete table in the main text reports aperture sums of these same difference images.

F1: detector-grid row 1, direct template reproduction

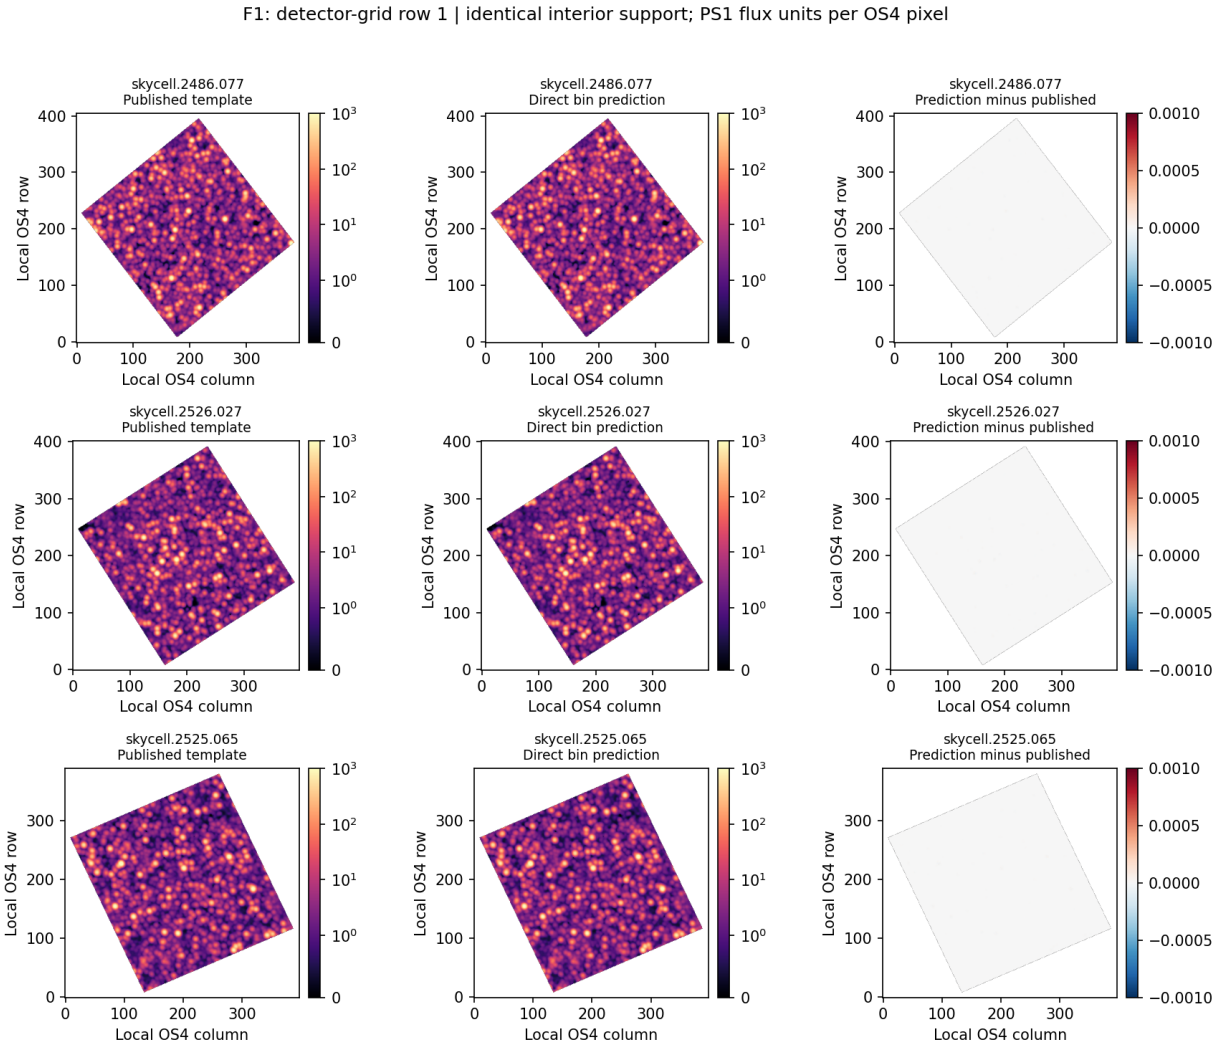


Figure 11. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2486.077	73,937	1.87×10^{-6}	0
skycell.2526.027	76,589	1.82×10^{-6}	0
skycell.2525.065	80,136	1.80×10^{-6}	0

F1: detector-grid row 2, direct template reproduction

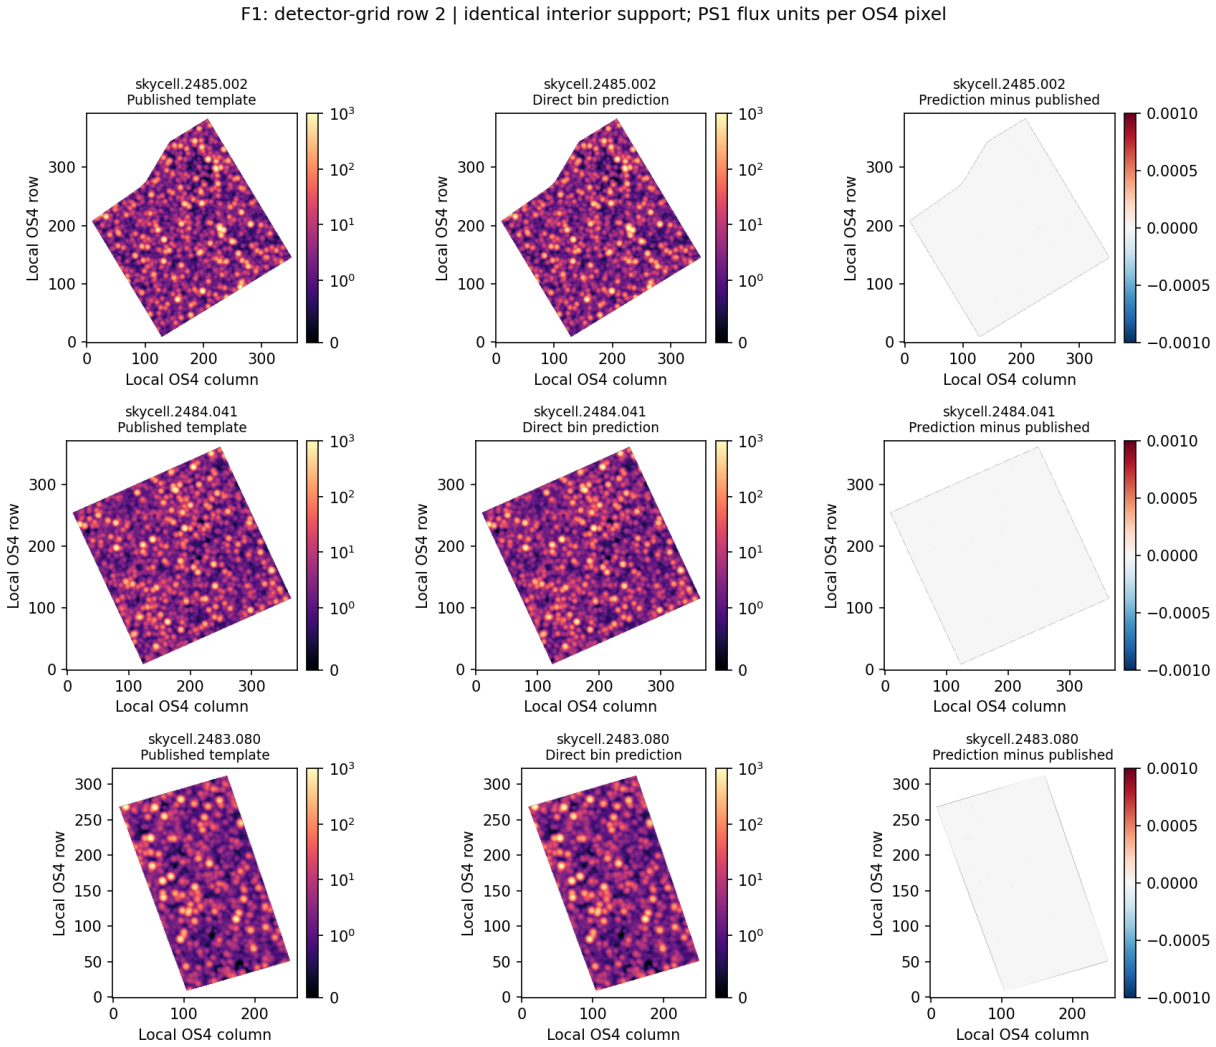


Figure 12. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2485.002	66,493	2.23×10^{-6}	0
skycell.2484.041	71,281	1.94×10^{-6}	0
skycell.2483.080	42,765	2.08×10^{-6}	0

F1: detector-grid row 3, direct template reproduction

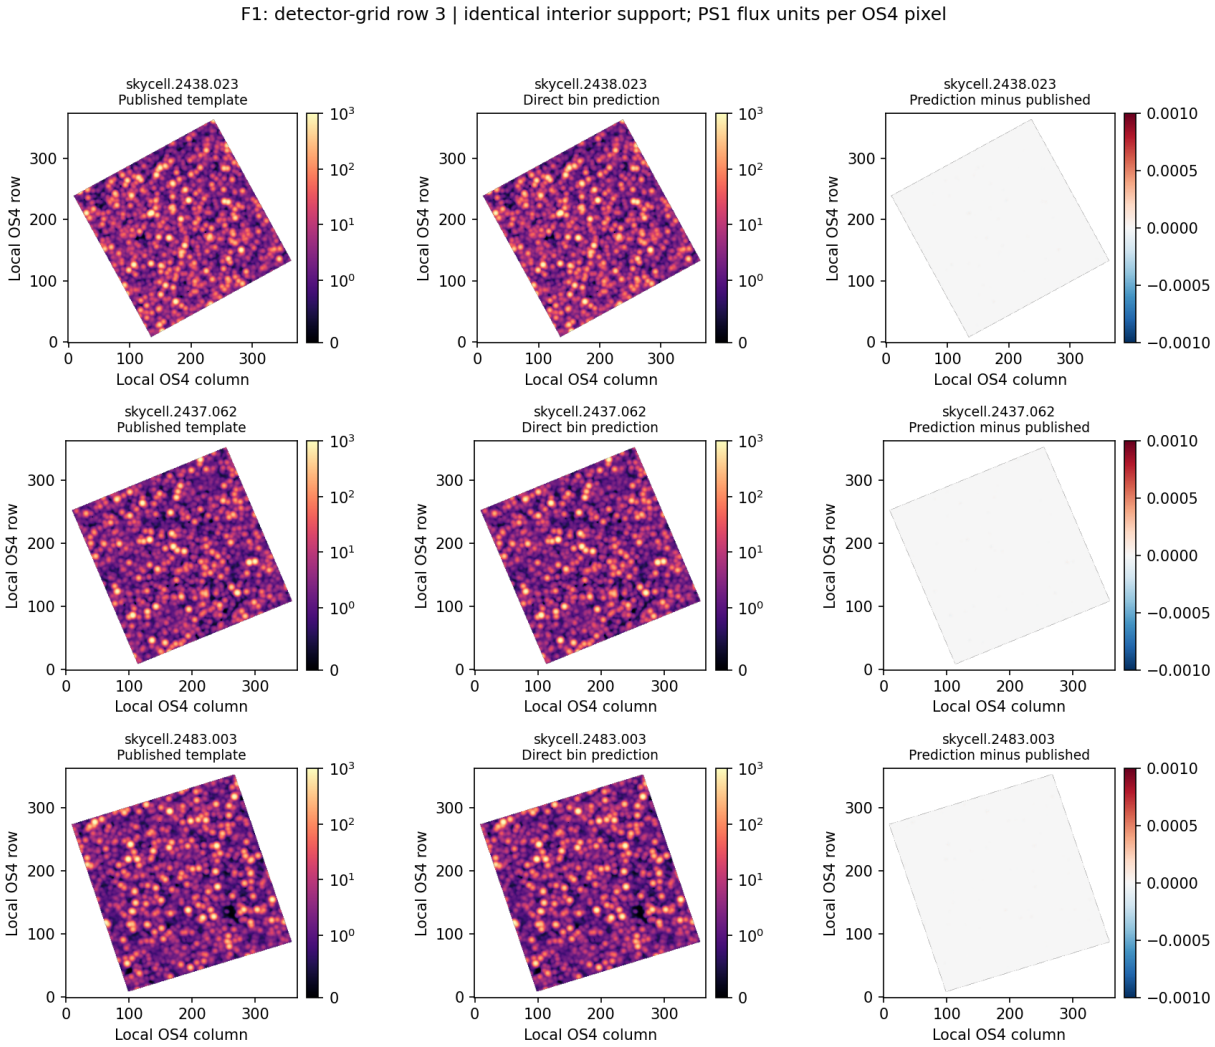


Figure 13. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2438.023	68,464	2.37×10^{-6}	0
skycell.2437.062	69,944	2.12×10^{-6}	0
skycell.2483.003	75,582	2.07×10^{-6}	0

F2: all four original PS1 bands and source removal

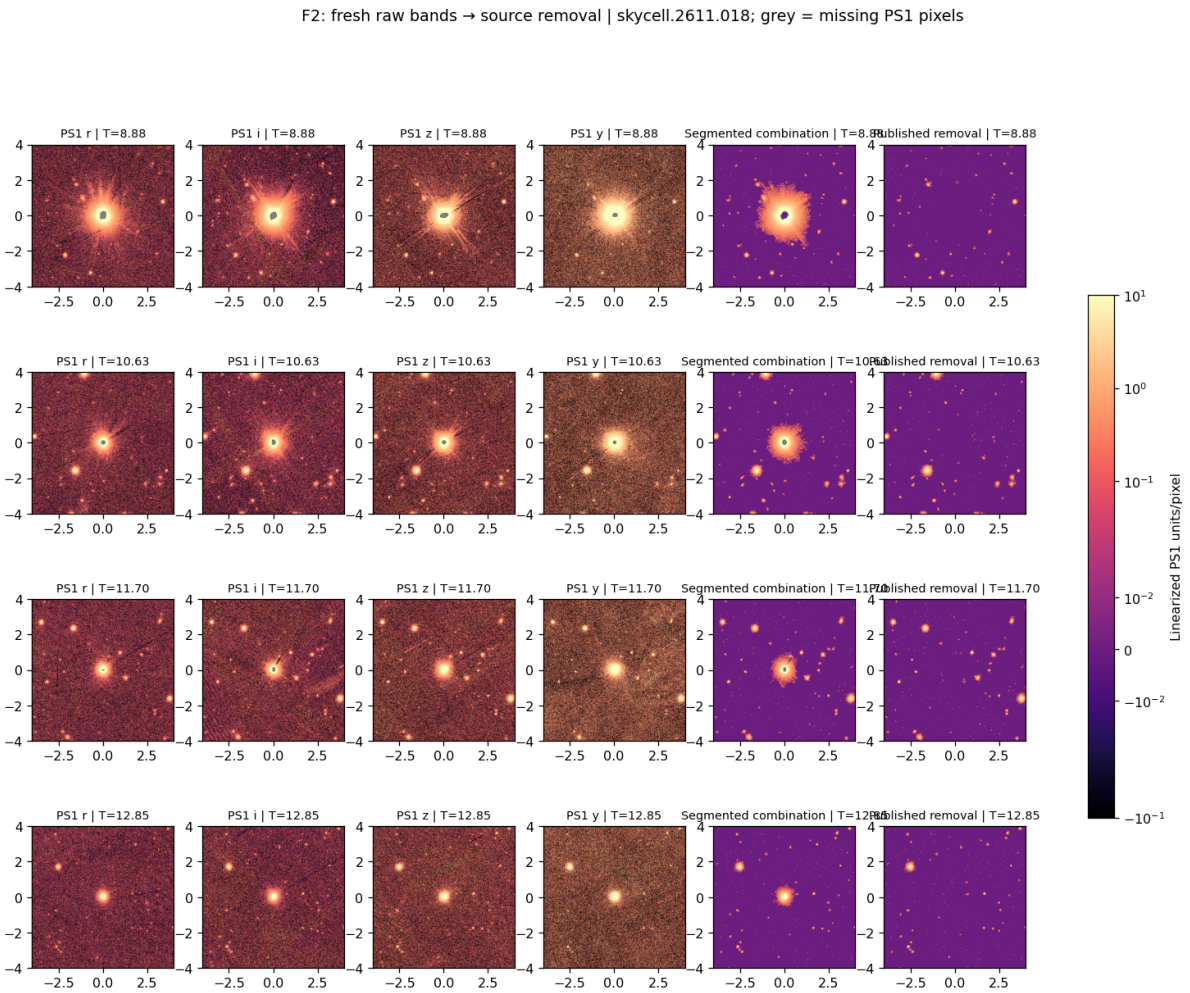


Figure 14. Four magnitude-spread examples, identical source positions. Columns show the original r, i, z, y bands, segmented combination and published removal. Grey marks unavailable original PS1 pixels; intensities in different bands differ physically. The centre-cell full-array replay is exact. The raw PS1 values have their own flux convention, not TESS electron/s.

F2: absolute stage images and paired differences

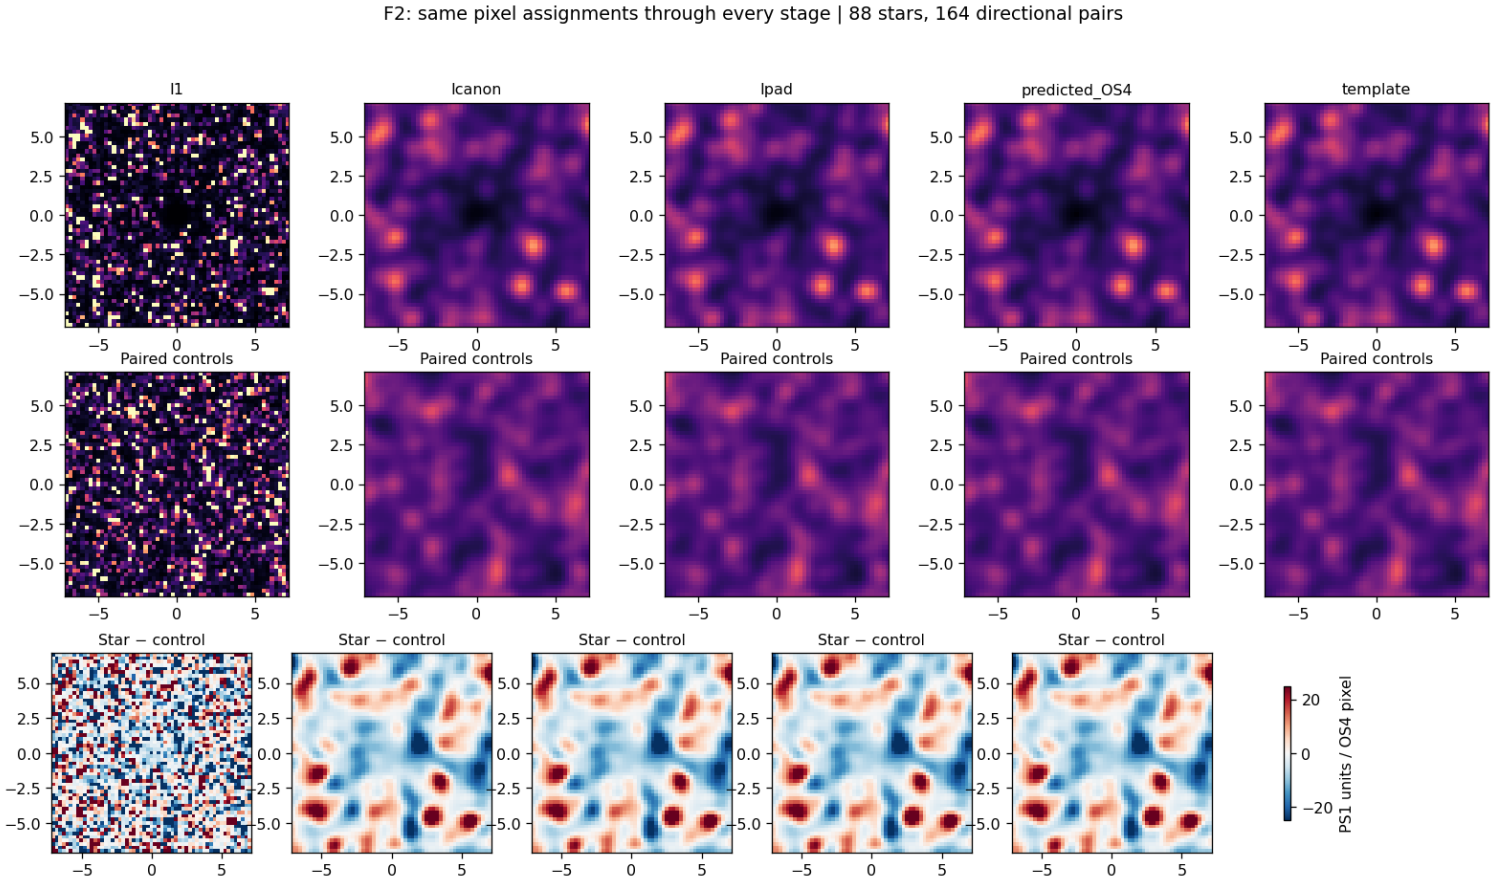


Figure 15. Top: mean at removed-star positions. Middle: paired directional controls. Bottom: their difference. Columns progress from published source-removed PS1 pixels through canonical and padding-aware convolution to predicted and published OS4 template. The central depression is already present in the first column; the later calculations propagate it. Native PS1 sums have been put on the same frozen OS4 assignments; units are PS1 flux per OS4 pixel.

F2: an individual patch and the mean saved-kernel replay

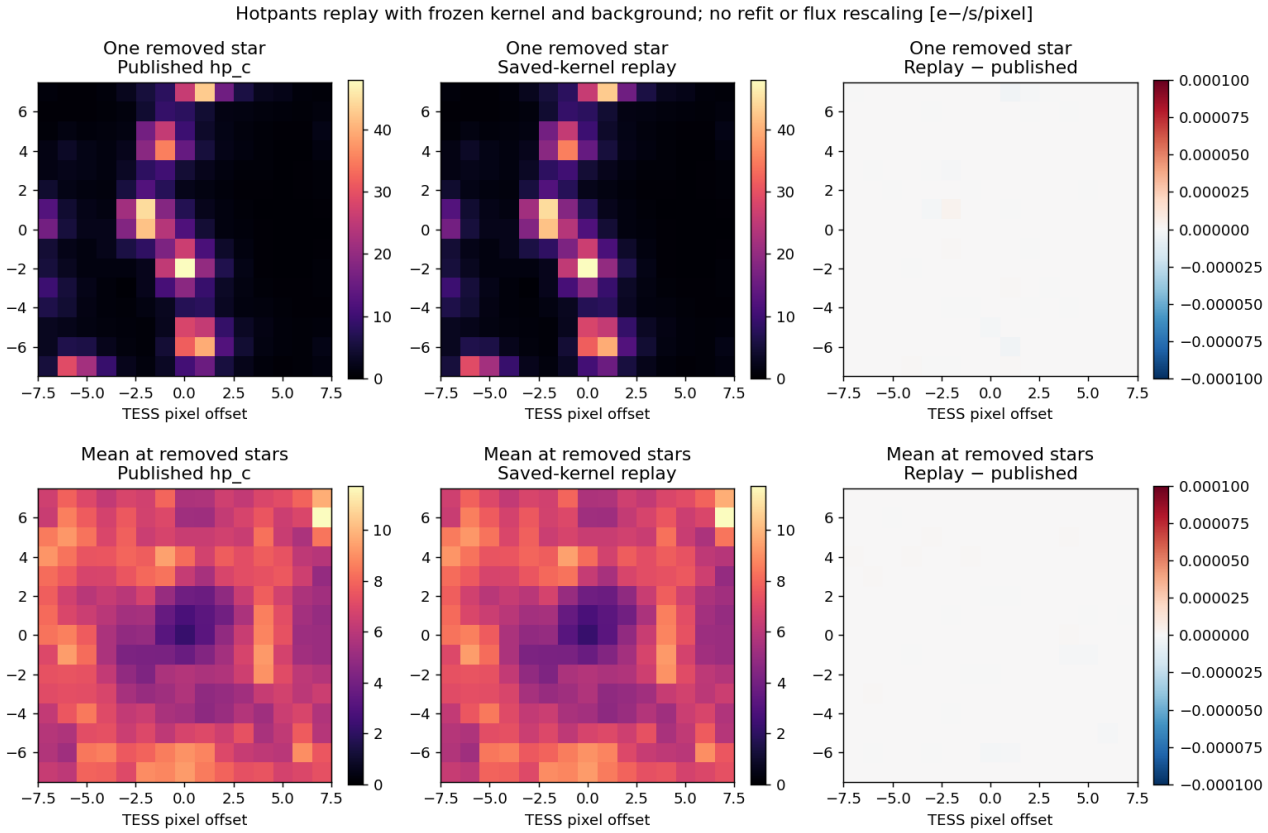


Figure 16. Published final template, replay with the frozen field kernel/background, and their difference. The first row is an individual removed-star patch; the second is the mean over tested star positions. The residual has its own narrow scale. No fit or flux rescaling was used; units are electron/s per native pixel.

F2: untouched-sky selection null

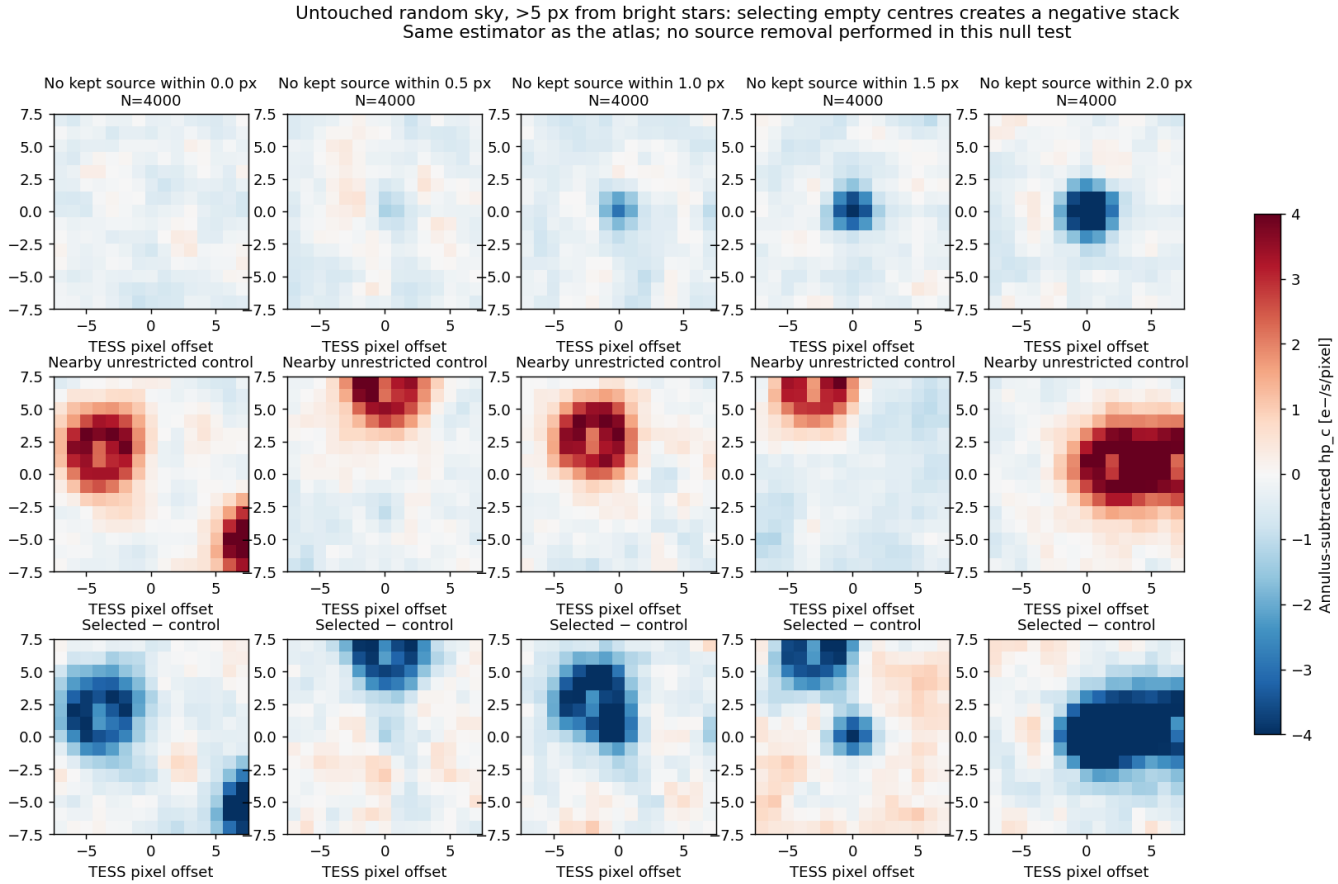


Figure 17. Columns change only the minimum distance to retained Gaia sources. Rows show selected unchanged locations, their nearby unrestricted controls and the difference. Both centres are away from bright stars. All images use the same scale in electron/s/native pixel after local annulus subtraction. A depression can be generated by changing the source population sampled by the estimator. Strong off-centre structures in the F2 controls account for its larger uncertainties.

F2: real stars, matched controls and both contrasts

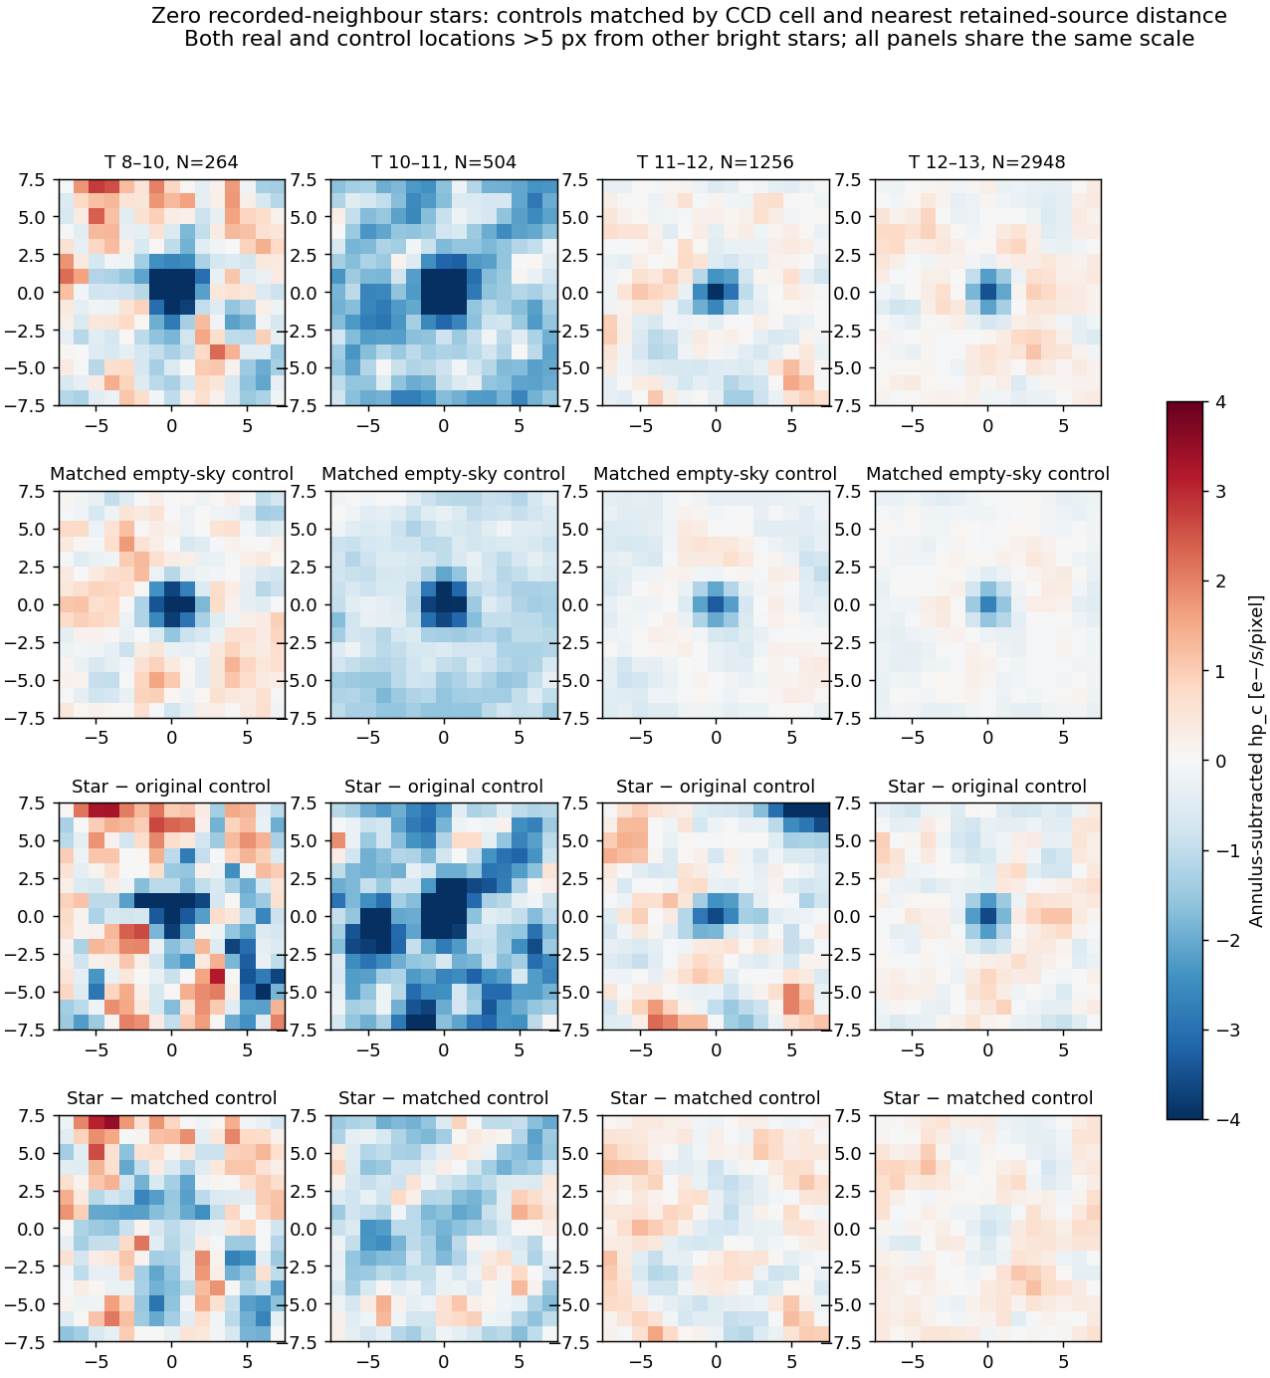


Figure 18. Columns are the four catalogue-magnitude bins. Rows show absolute annulus-subtracted star stamps, matched controls, the original contrast and the matched contrast. Matching never uses central template intensity. All panels share one scale, in electron/s/native pixel. The complete table in the main text reports aperture sums of these same difference images.

F2: detector-grid row 1, direct template reproduction

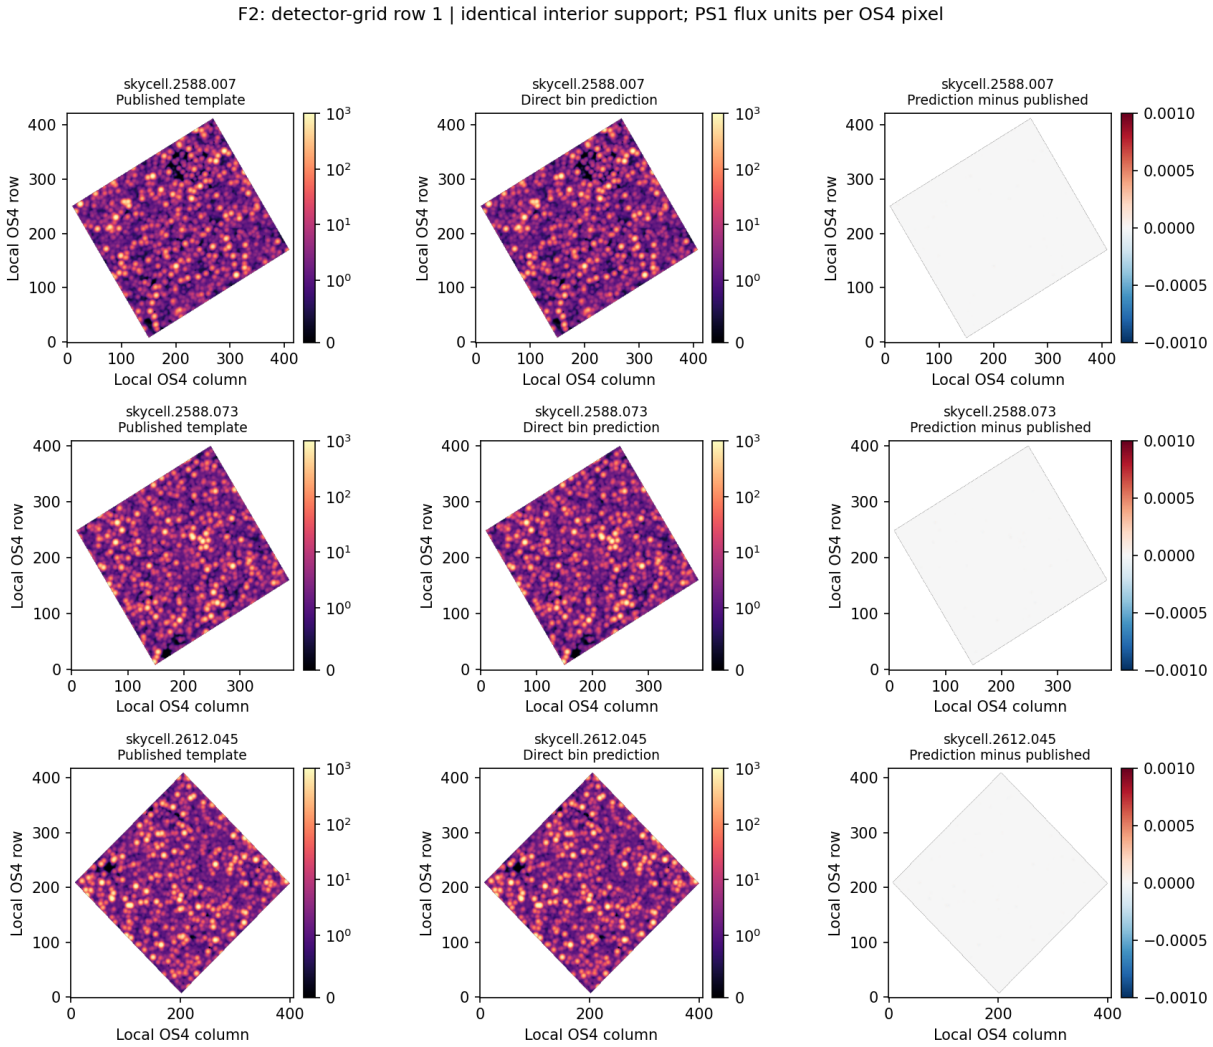


Figure 19. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2588.007	85,457	1.69×10^{-6}	0
skycell.2588.073	78,534	1.91×10^{-6}	0
skycell.2612.045	78,709	2.44×10^{-6}	0

F2: detector-grid row 2, direct template reproduction

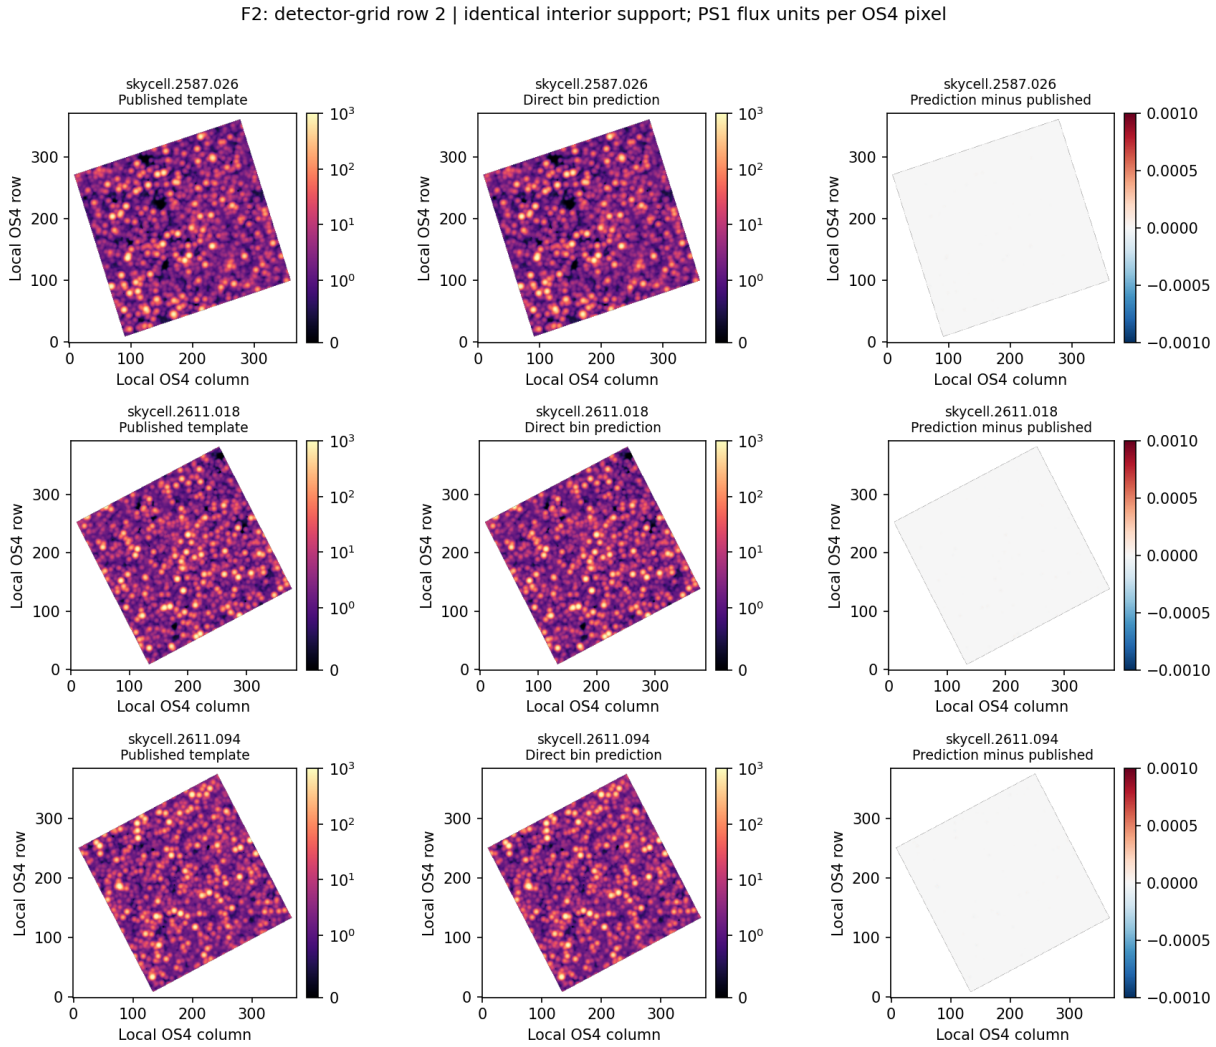


Figure 20. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2587.026	78,108	1.81×10^{-6}	0
skycell.2611.018	75,529	2.19×10^{-6}	0
skycell.2611.094	71,838	2.32×10^{-6}	0

F2: detector-grid row 3, direct template reproduction

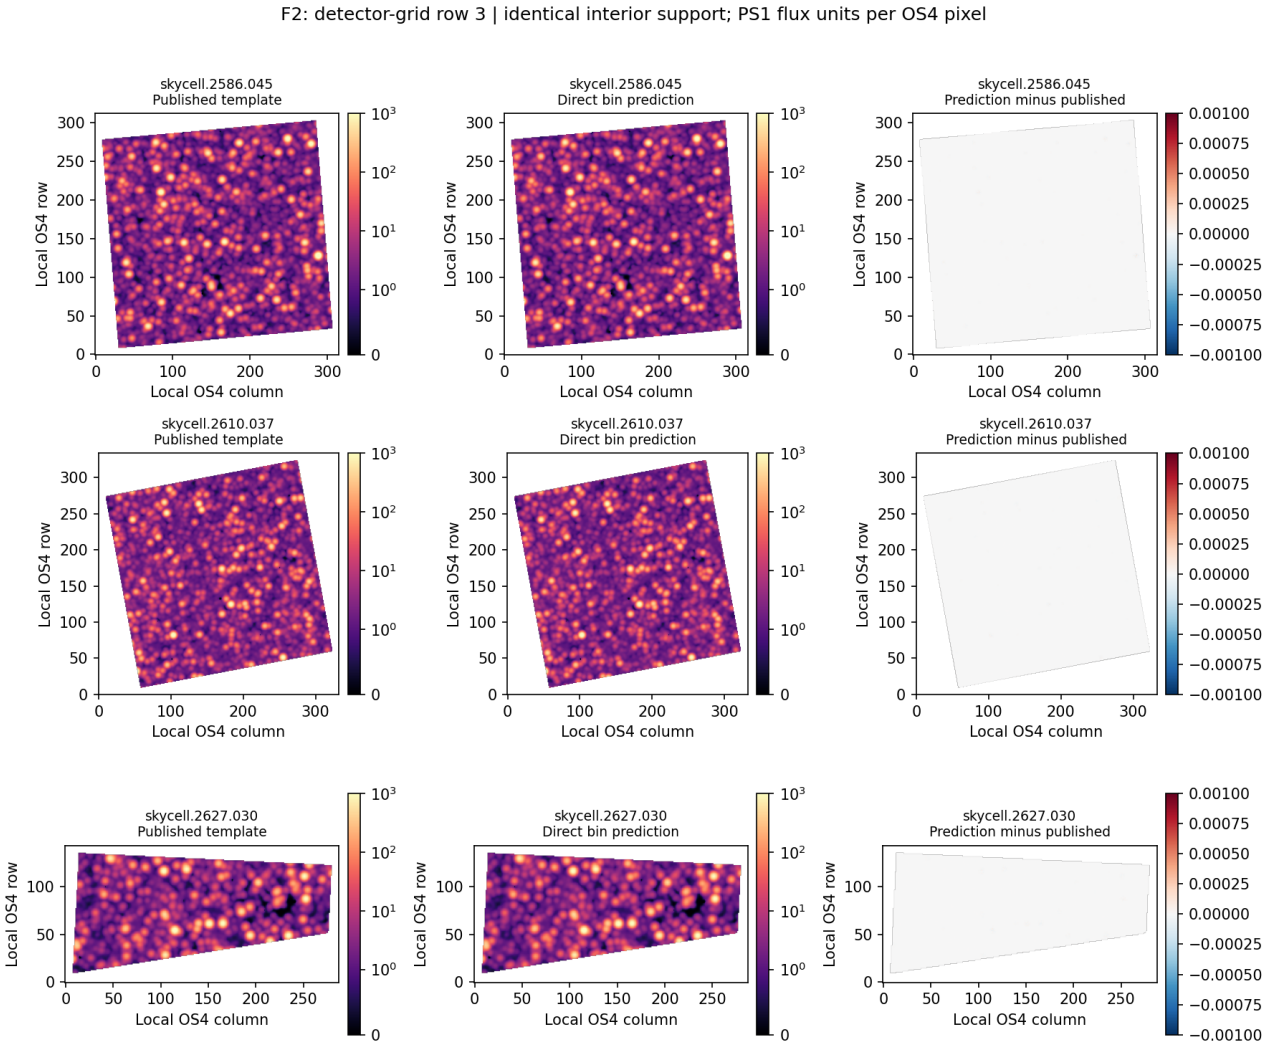


Figure 21. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2586.045	75,413	2.06×10^{-6}	0
skycell.2610.037	72,801	1.49×10^{-6}	0
skycell.2627.030	26,307	2.24×10^{-6}	0

S22: all four original PS1 bands and source removal

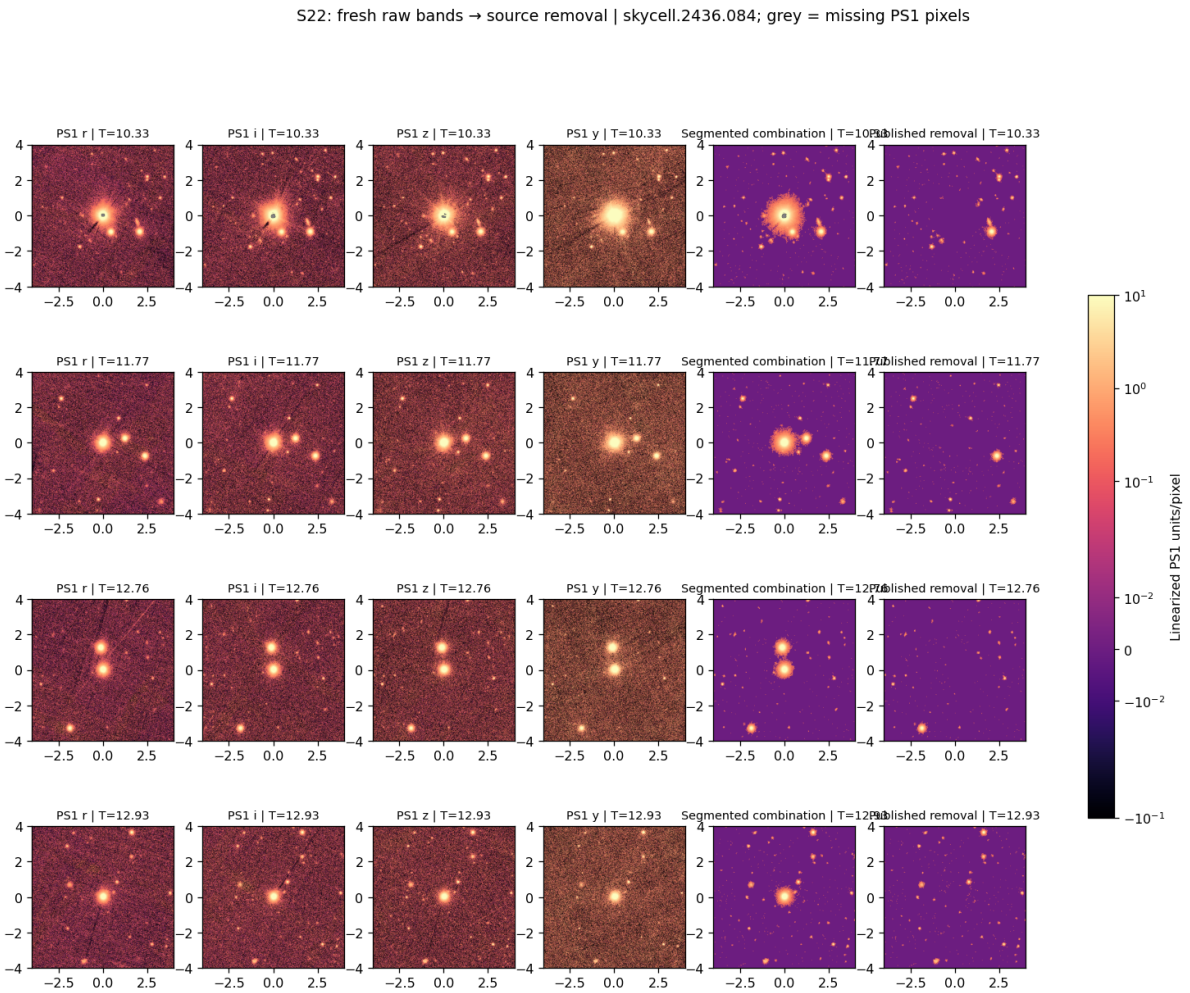


Figure 22. Four magnitude-spread examples, identical source positions. Columns show the original r, i, z, y bands, segmented combination and published removal. Grey marks unavailable original PS1 pixels; intensities in different bands differ physically. The centre-cell full-array replay is exact. The raw PS1 values have their own flux convention, not TESS electron/s.

S22: absolute stage images and paired differences

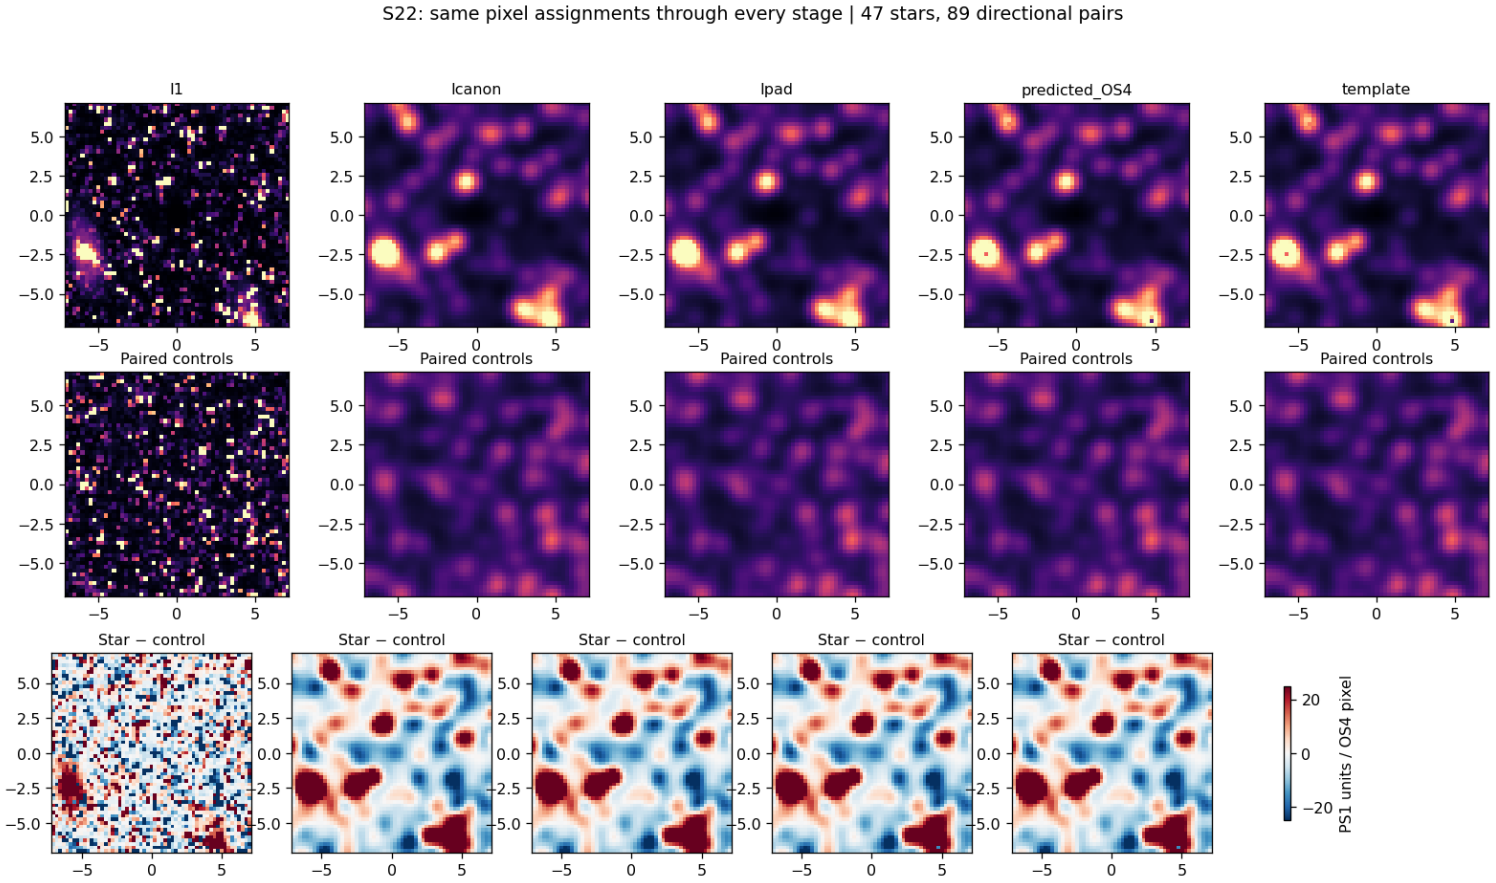


Figure 23. Top: mean at removed-star positions. Middle: paired directional controls. Bottom: their difference. Columns progress from published source-removed PS1 pixels through canonical and padding-aware convolution to predicted and published OS4 template. The central depression is already present in the first column; the later calculations propagate it. Native PS1 sums have been put on the same frozen OS4 assignments; units are PS1 flux per OS4 pixel.

S22: an individual patch and the mean saved-kernel replay

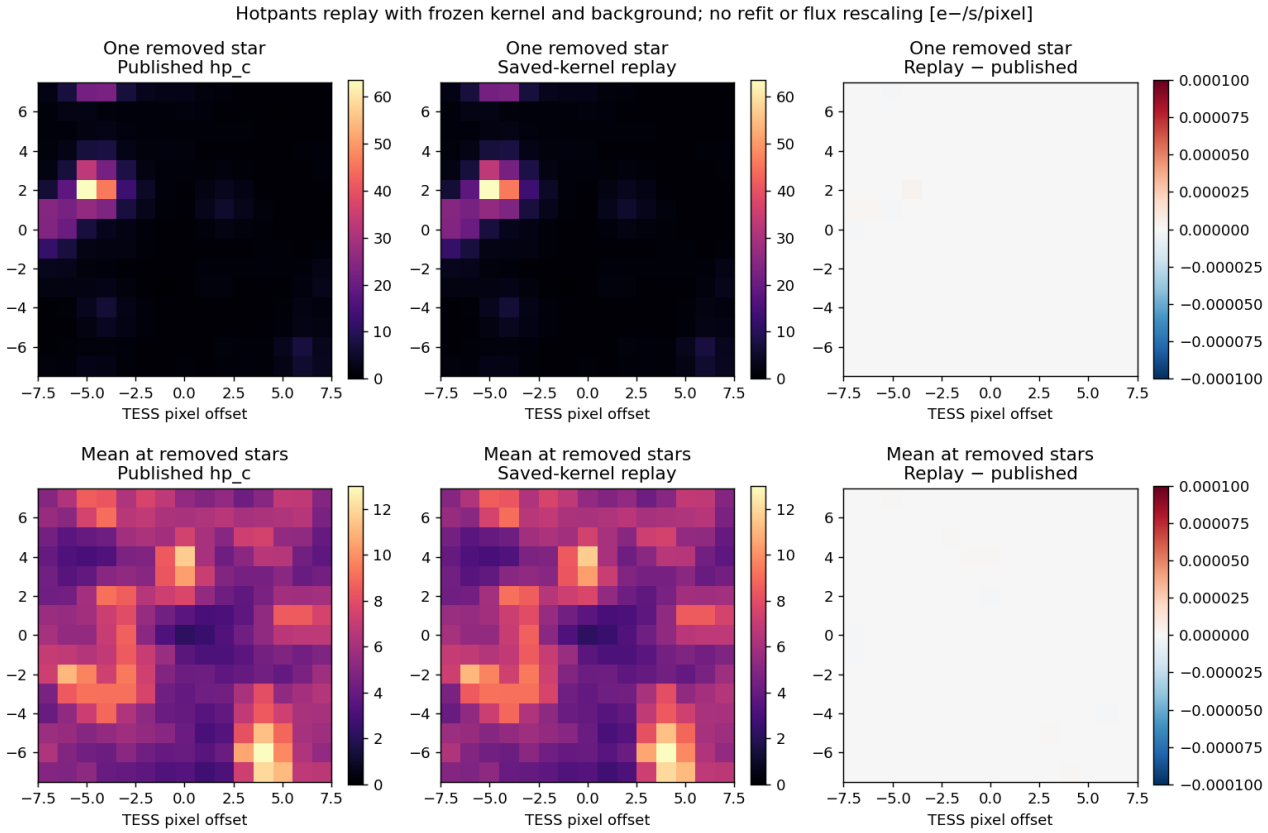


Figure 24. Published final template, replay with the frozen field kernel/background, and their difference. The first row is an individual removed-star patch; the second is the mean over tested star positions. The residual has its own narrow scale. No fit or flux rescaling was used; units are electron/s per native pixel.

S22: untouched-sky selection null

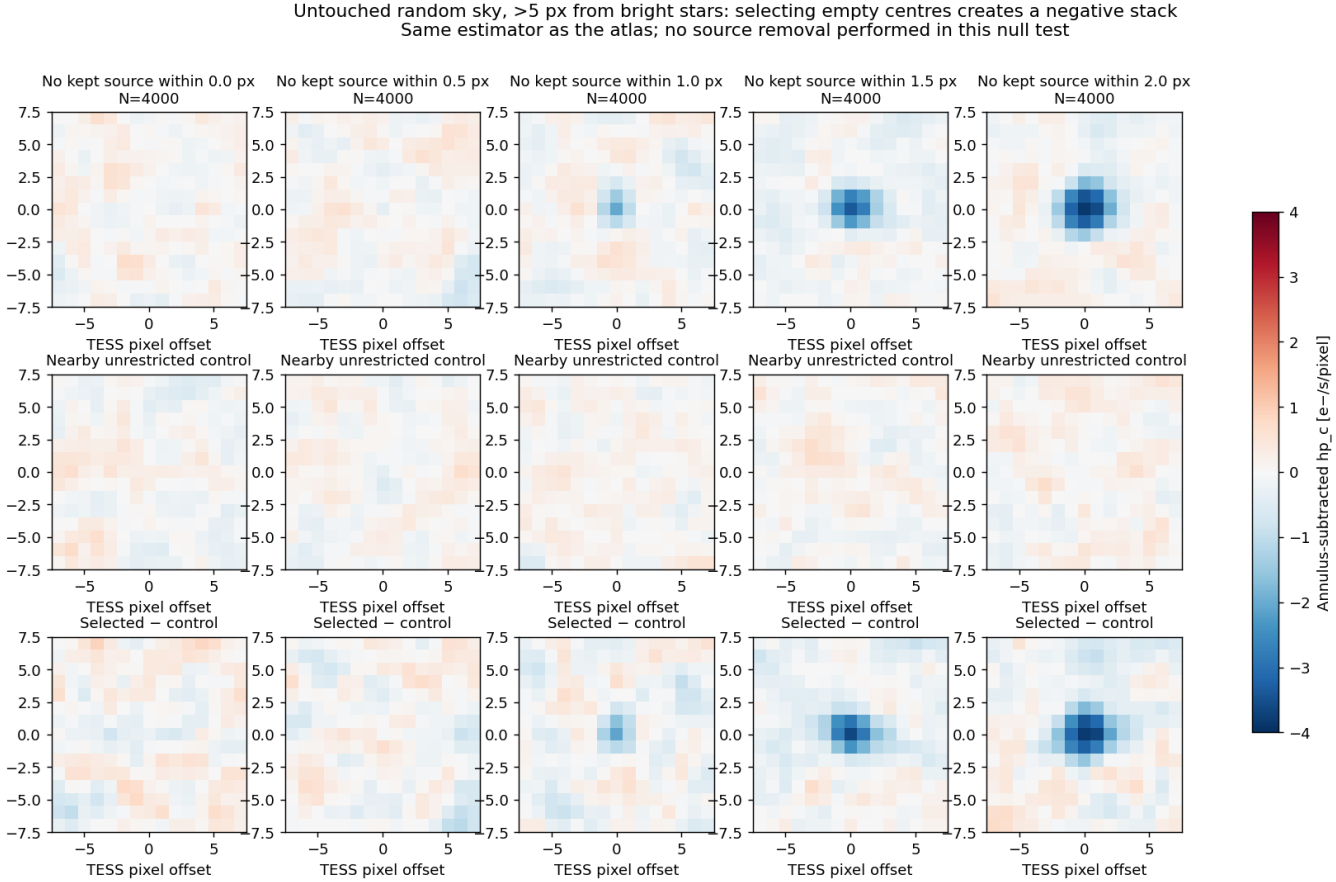


Figure 25. Columns change only the minimum distance to retained Gaia sources. Rows show selected unchanged locations, their nearby unrestricted controls and the difference. Both centres are away from bright stars. All images use the same scale in electron/s/native pixel after local annulus subtraction. A depression can be generated by changing the source population sampled by the estimator. Strong off-centre structures in the F2 controls account for its larger uncertainties.

S22: real stars, matched controls and both contrasts

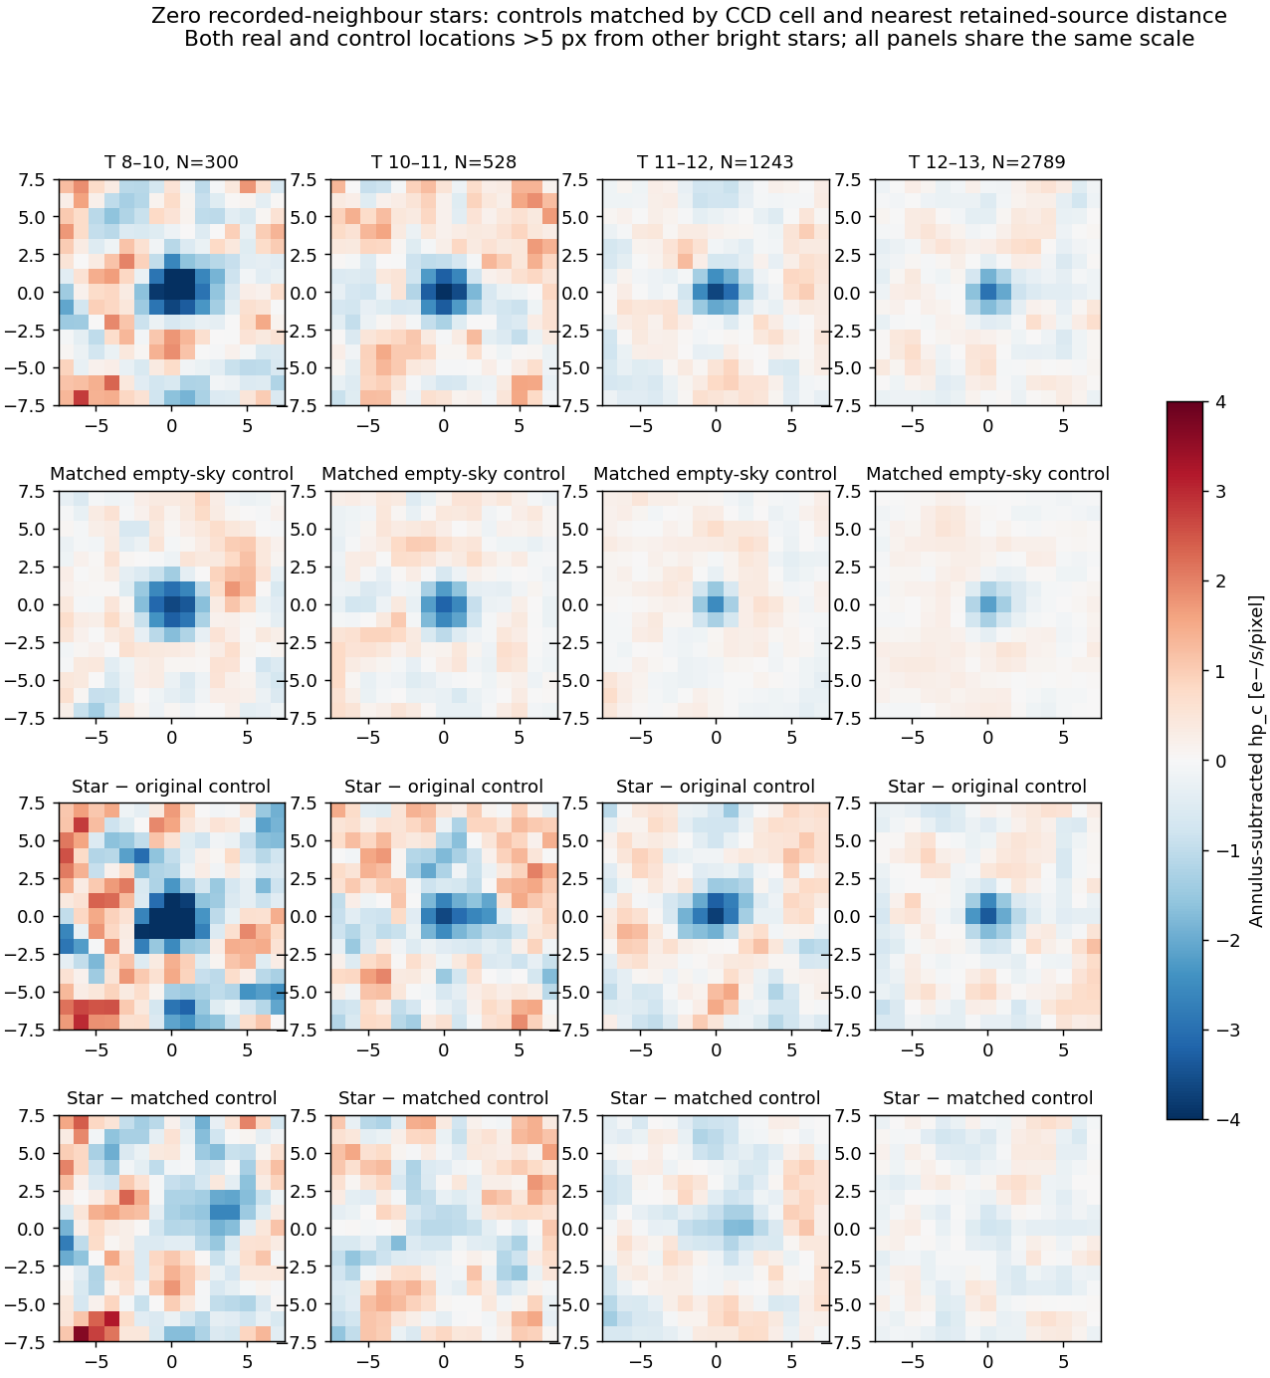


Figure 26. Columns are the four catalogue-magnitude bins. Rows show absolute annulus-subtracted star stamps, matched controls, the original contrast and the matched contrast. Matching never uses central template intensity. All panels share one scale, in electron/s/native pixel. The complete table in the main text reports aperture sums of these same difference images.

S22: detector-grid row 1, direct template reproduction

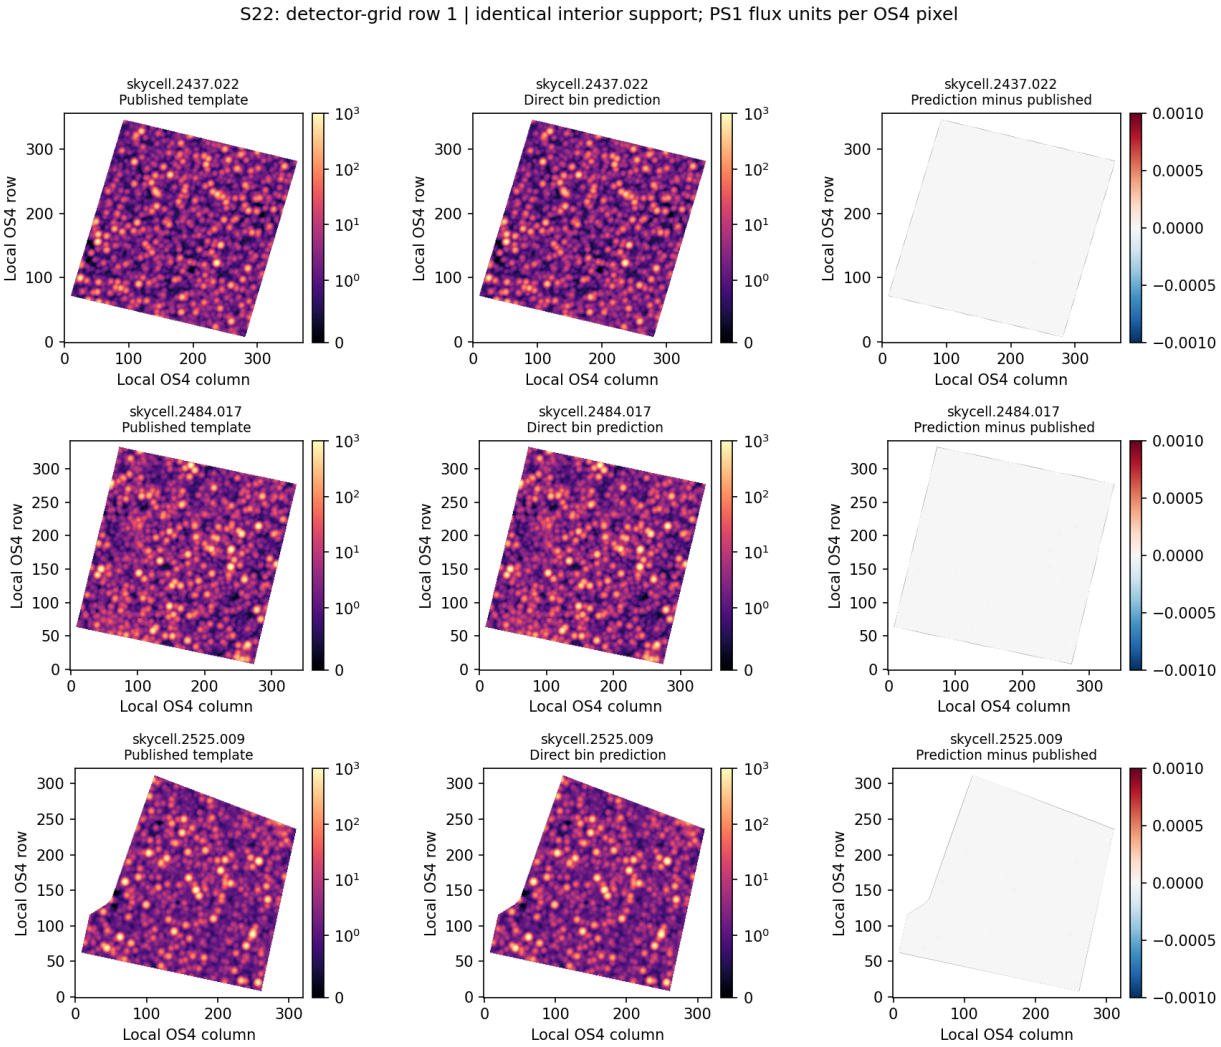


Figure 27. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2437.022	79,081	1.14×10^{-6}	0
skycell.2484.017	74,481	1.90×10^{-6}	0
skycell.2525.009	58,030	1.98×10^{-6}	0

S22: detector-grid row 2, direct template reproduction

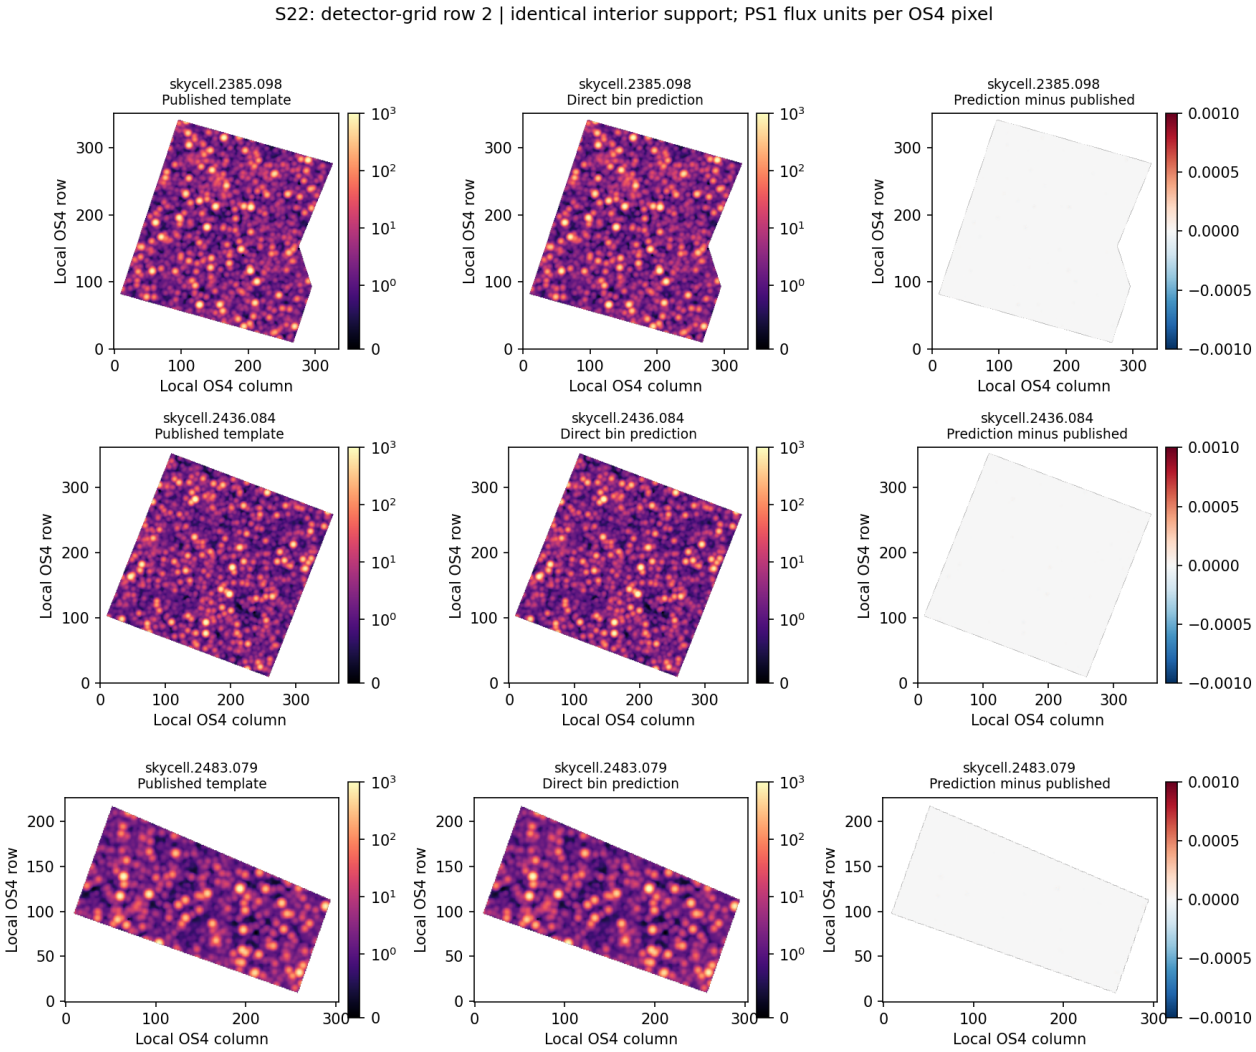


Figure 28. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2385.098	67,929	2.20×10^{-6}	0
skycell.2436.084	71,508	1.82×10^{-6}	0
skycell.2483.079	31,015	1.71×10^{-6}	0

S22: detector-grid row 3, direct template reproduction

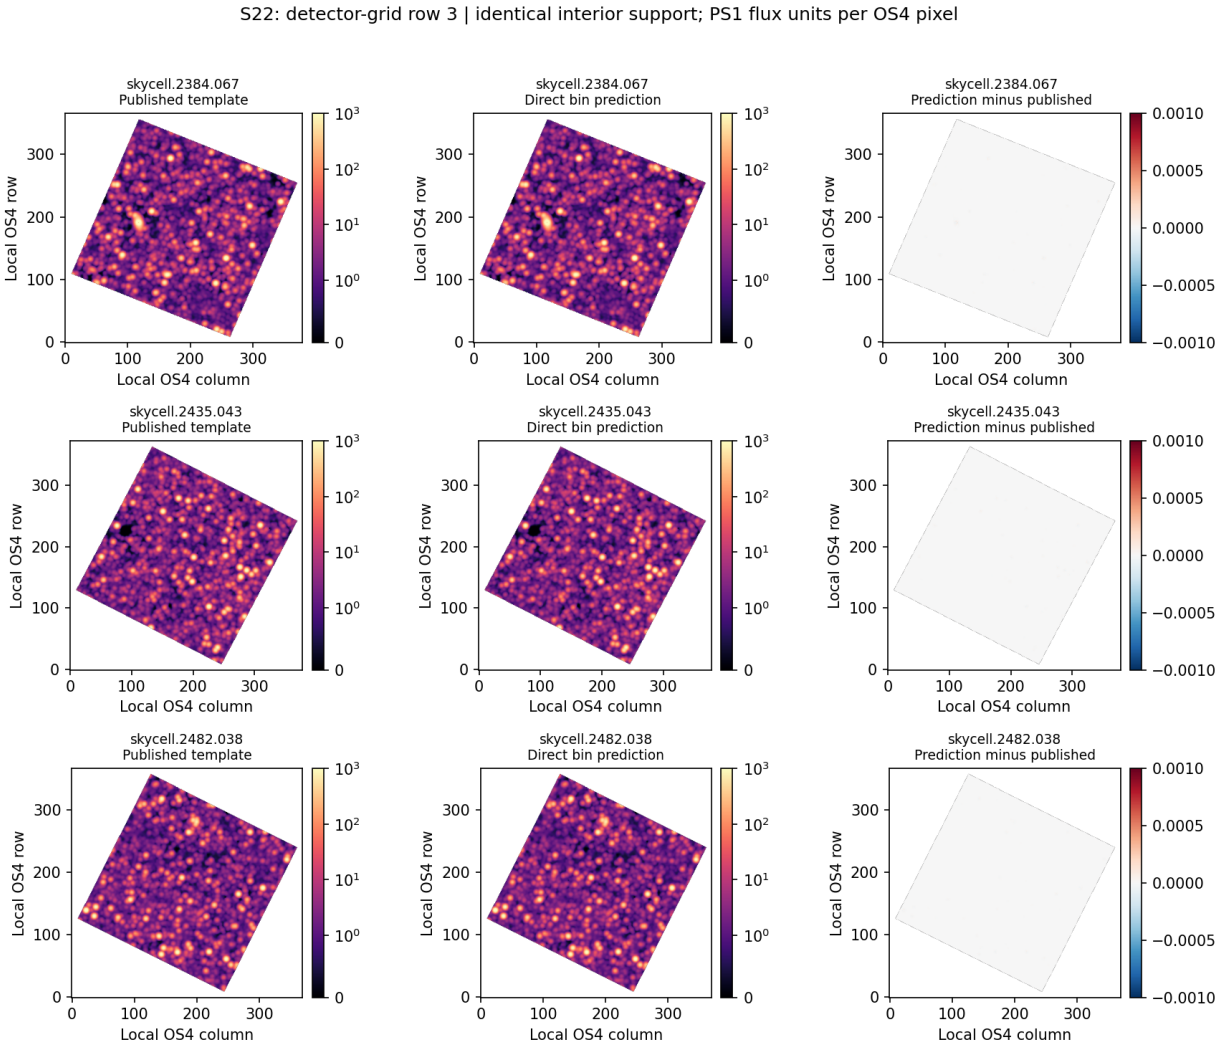


Figure 29. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2384.067	73,347	2.38×10^{-6}	0
skycell.2435.043	70,196	1.98×10^{-6}	0
skycell.2482.038	67,867	1.86×10^{-6}	0

C1: all four original PS1 bands and source removal

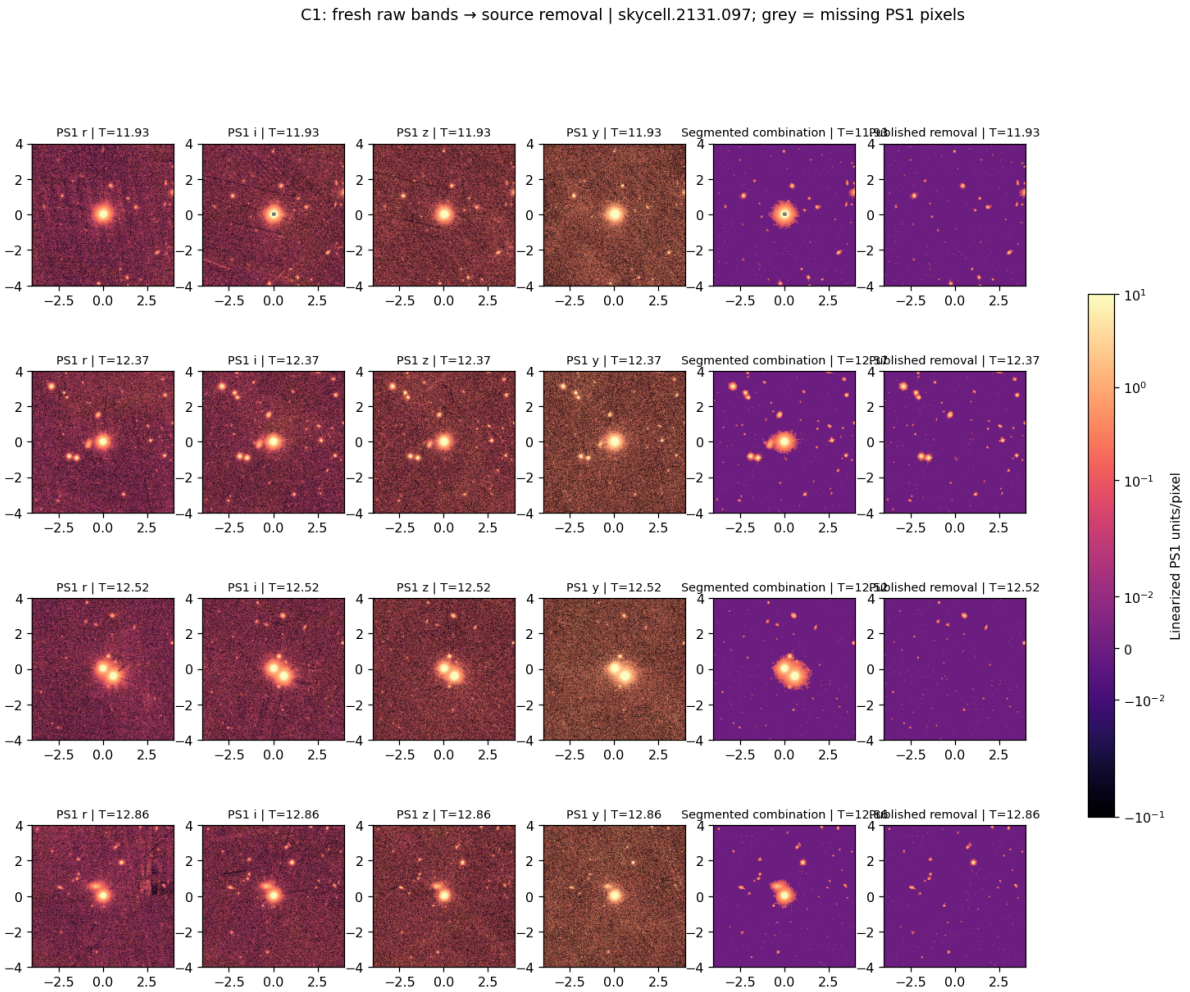


Figure 30. Four magnitude-spread examples, identical source positions. Columns show the original r, i, z, y bands, segmented combination and published removal. Grey marks unavailable original PS1 pixels; intensities in different bands differ physically. The centre-cell full-array replay is exact. The raw PS1 values have their own flux convention, not TESS electron/s.

C1: absolute stage images and paired differences

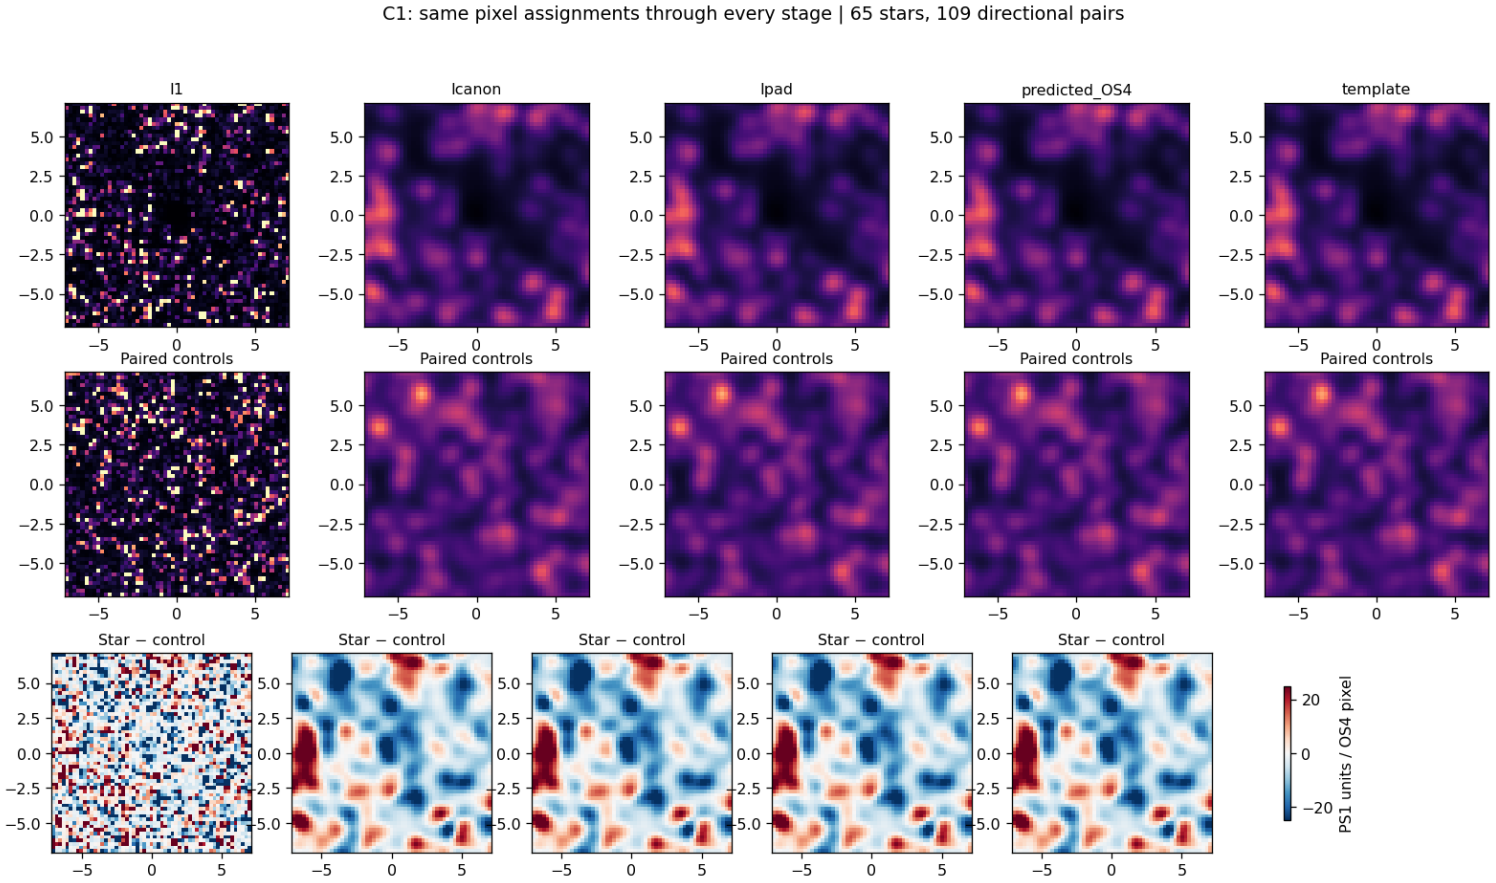


Figure 31. Top: mean at removed-star positions. Middle: paired directional controls. Bottom: their difference. Columns progress from published source-removed PS1 pixels through canonical and padding-aware convolution to predicted and published OS4 template. The central depression is already present in the first column; the later calculations propagate it. Native PS1 sums have been put on the same frozen OS4 assignments; units are PS1 flux per OS4 pixel.

C1: an individual patch and the mean saved-kernel replay

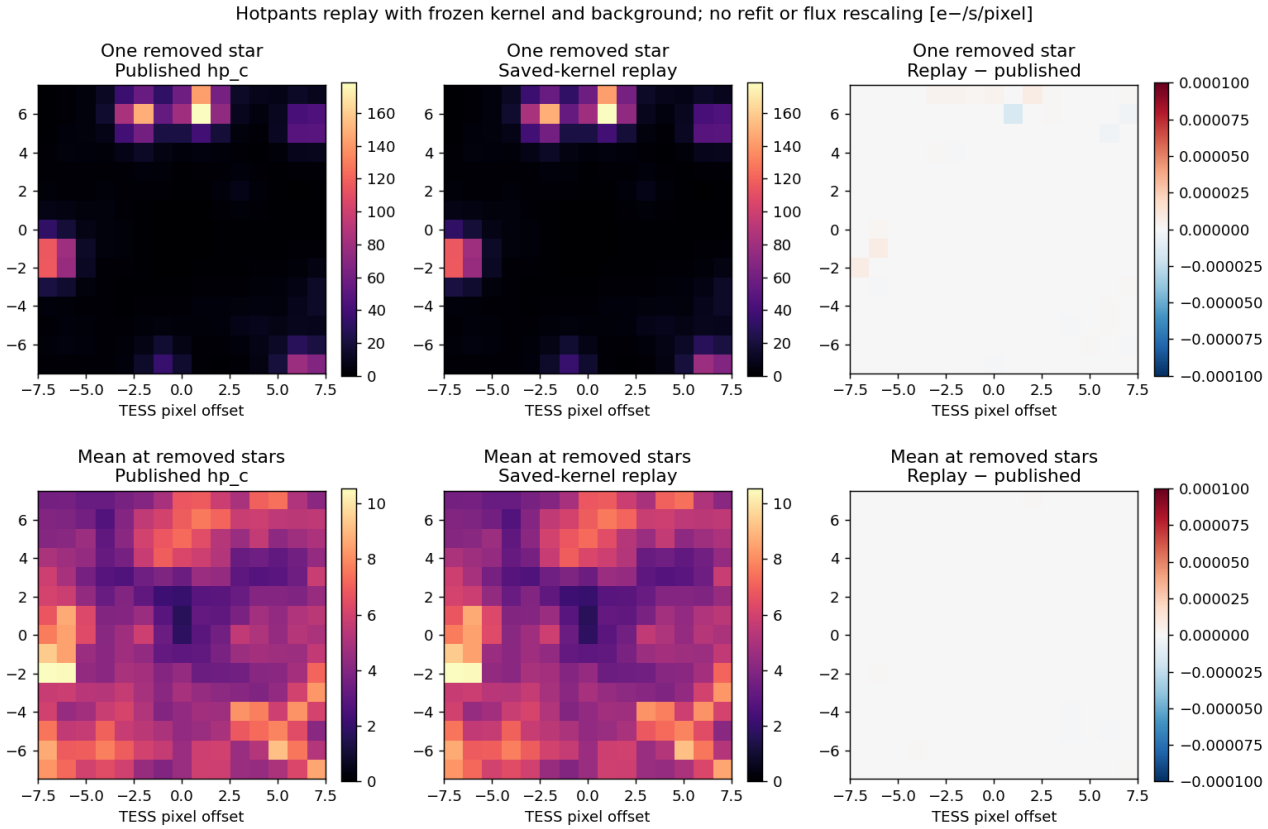


Figure 32. Published final template, replay with the frozen field kernel/background, and their difference. The first row is an individual removed-star patch; the second is the mean over tested star positions. The residual has its own narrow scale. No fit or flux rescaling was used; units are electron/s per native pixel.

C1: untouched-sky selection null

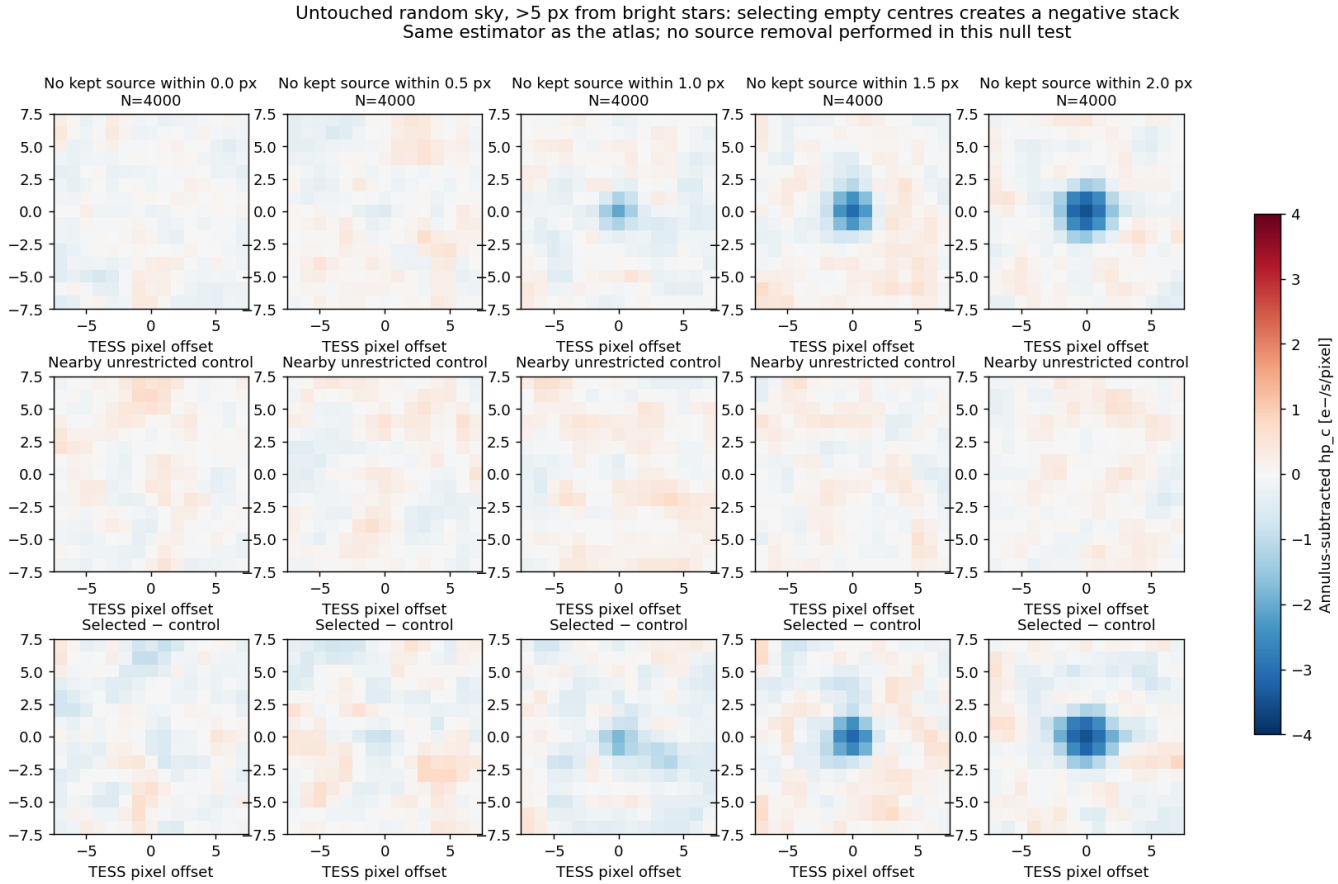


Figure 33. Columns change only the minimum distance to retained Gaia sources. Rows show selected unchanged locations, their nearby unrestricted controls and the difference. Both centres are away from bright stars. All images use the same scale in electron/s/native pixel after local annulus subtraction. A depression can be generated by changing the source population sampled by the estimator. Strong off-centre structures in the F2 controls account for its larger uncertainties.

C1: real stars, matched controls and both contrasts

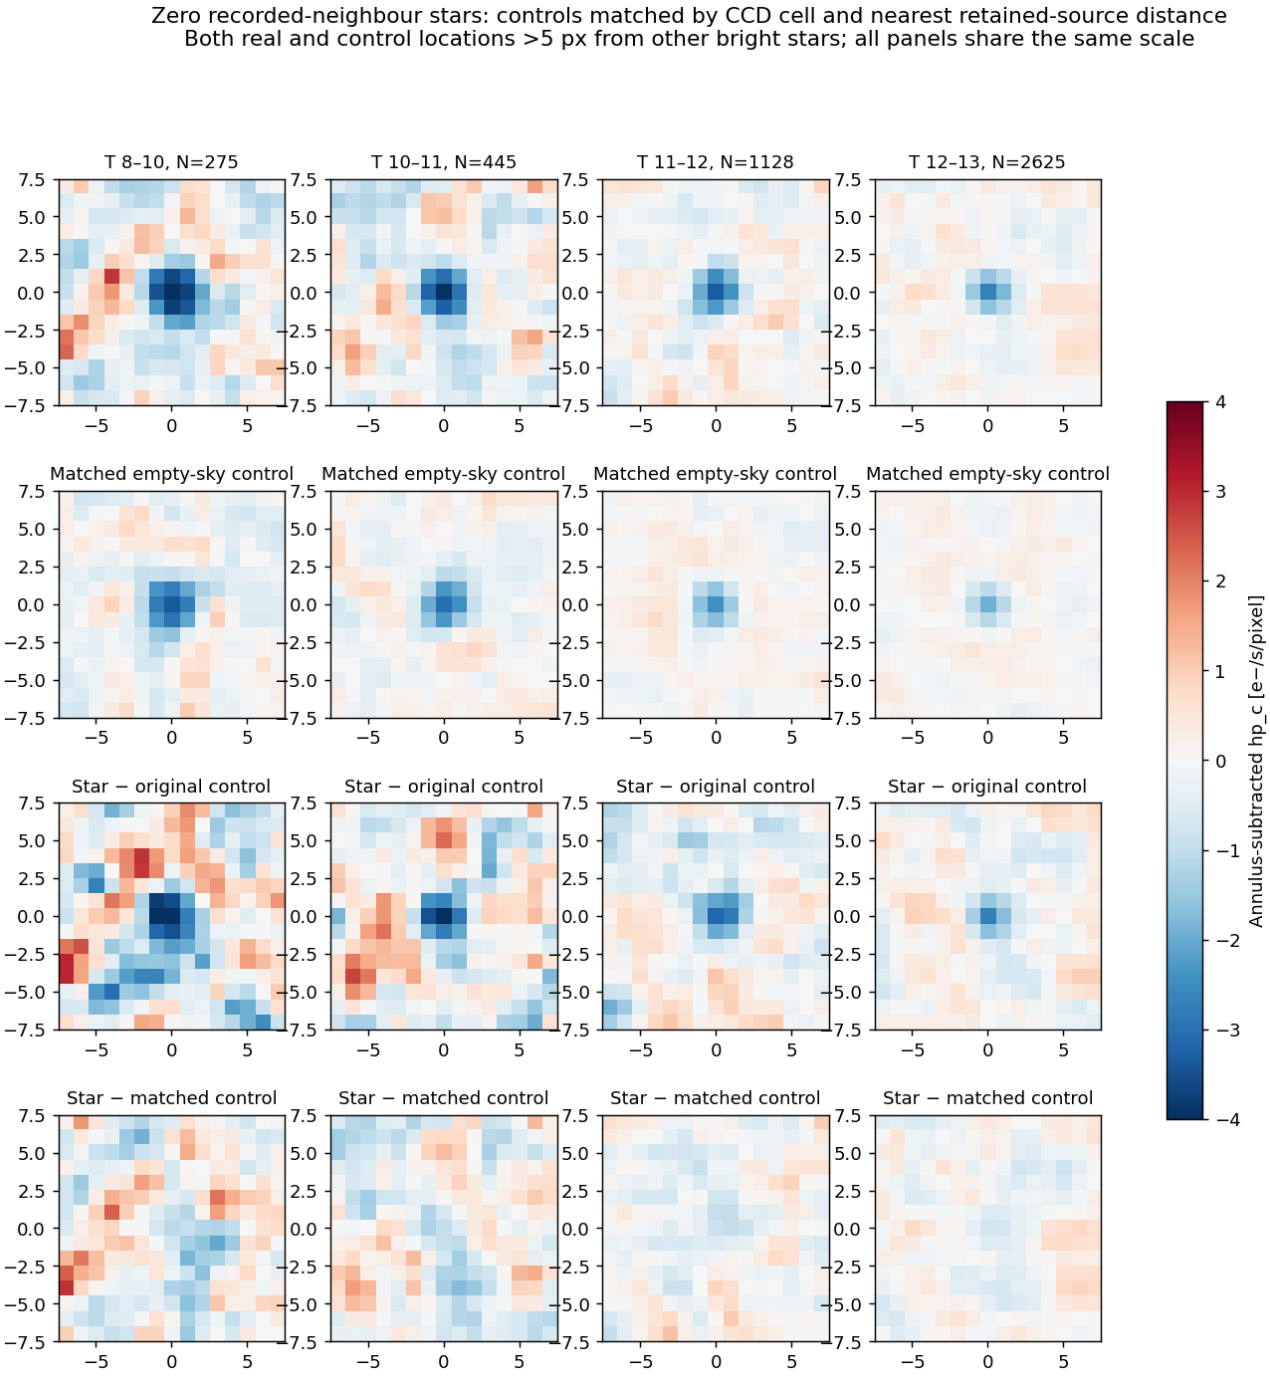


Figure 34. Columns are the four catalogue-magnitude bins. Rows show absolute annulus-subtracted star stamps, matched controls, the original contrast and the matched contrast. Matching never uses central template intensity. All panels share one scale, in electron/s/native pixel. The complete table in the main text reports aperture sums of these same difference images.

C1: detector-grid row 1, direct template reproduction

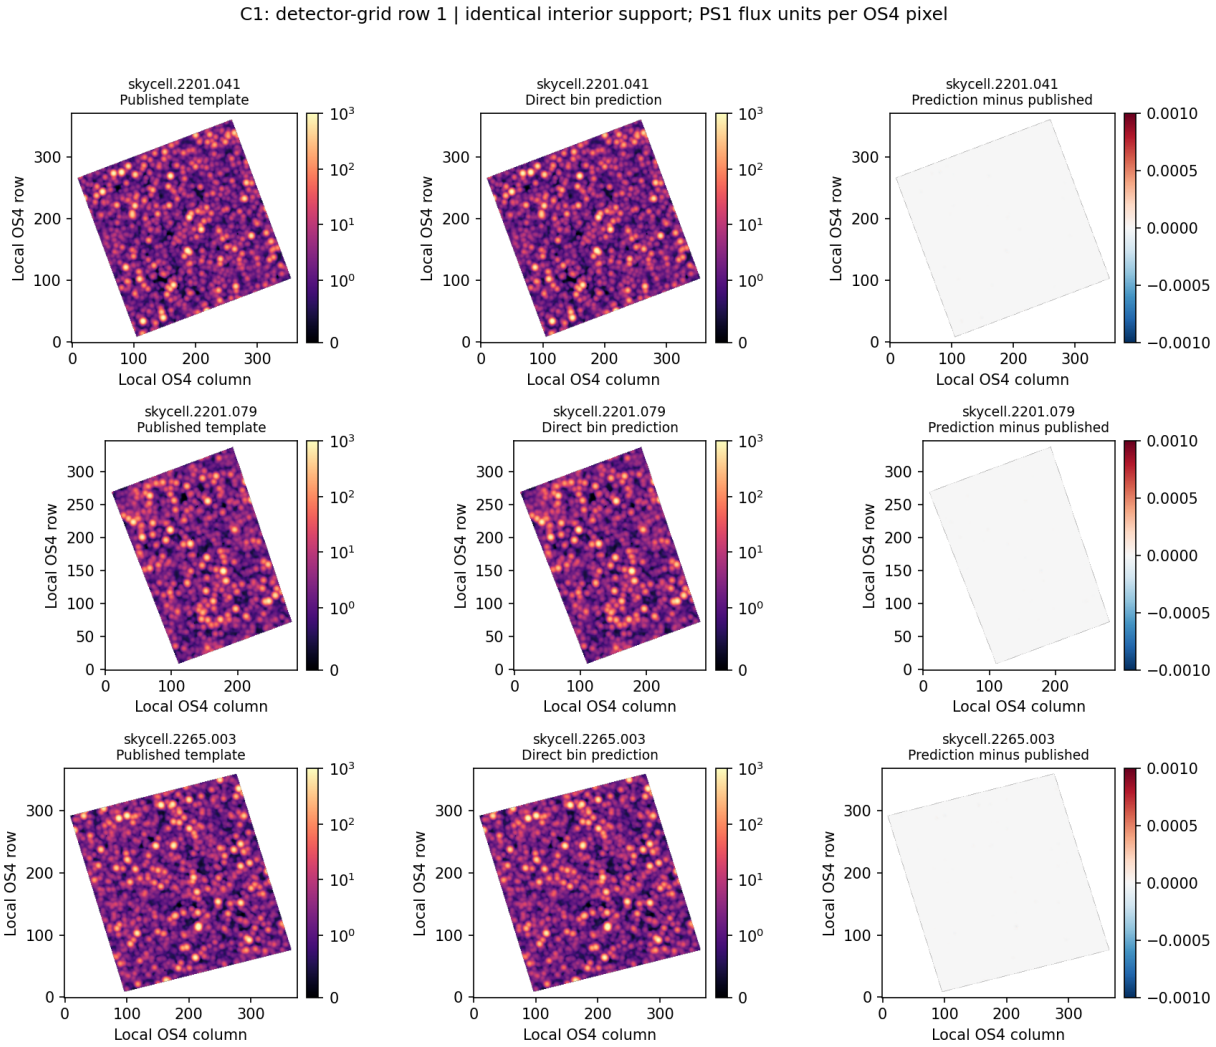


Figure 35. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2201.041	73,279	1.67×10^{-6}	0
skycell.2201.079	52,684	1.48×10^{-6}	0
skycell.2265.003	82,068	2.00×10^{-6}	0

C1: detector-grid row 2, direct template reproduction

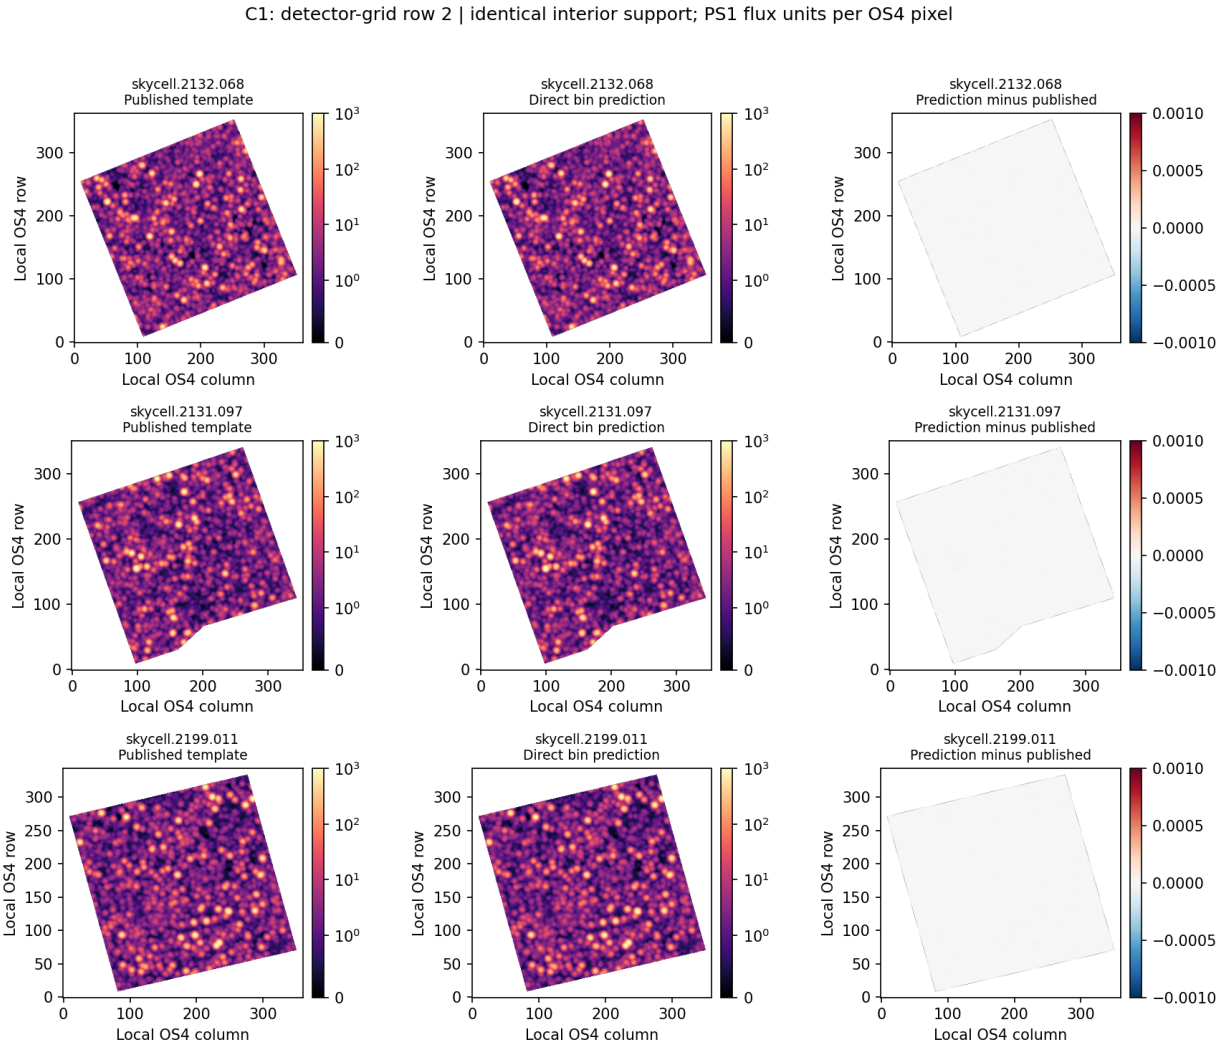


Figure 36. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2132.068	69,793	1.71×10^{-6}	0
skycell.2131.097	66,411	1.76×10^{-6}	0
skycell.2199.011	75,011	1.78×10^{-6}	0

C1: detector-grid row 3, direct template reproduction

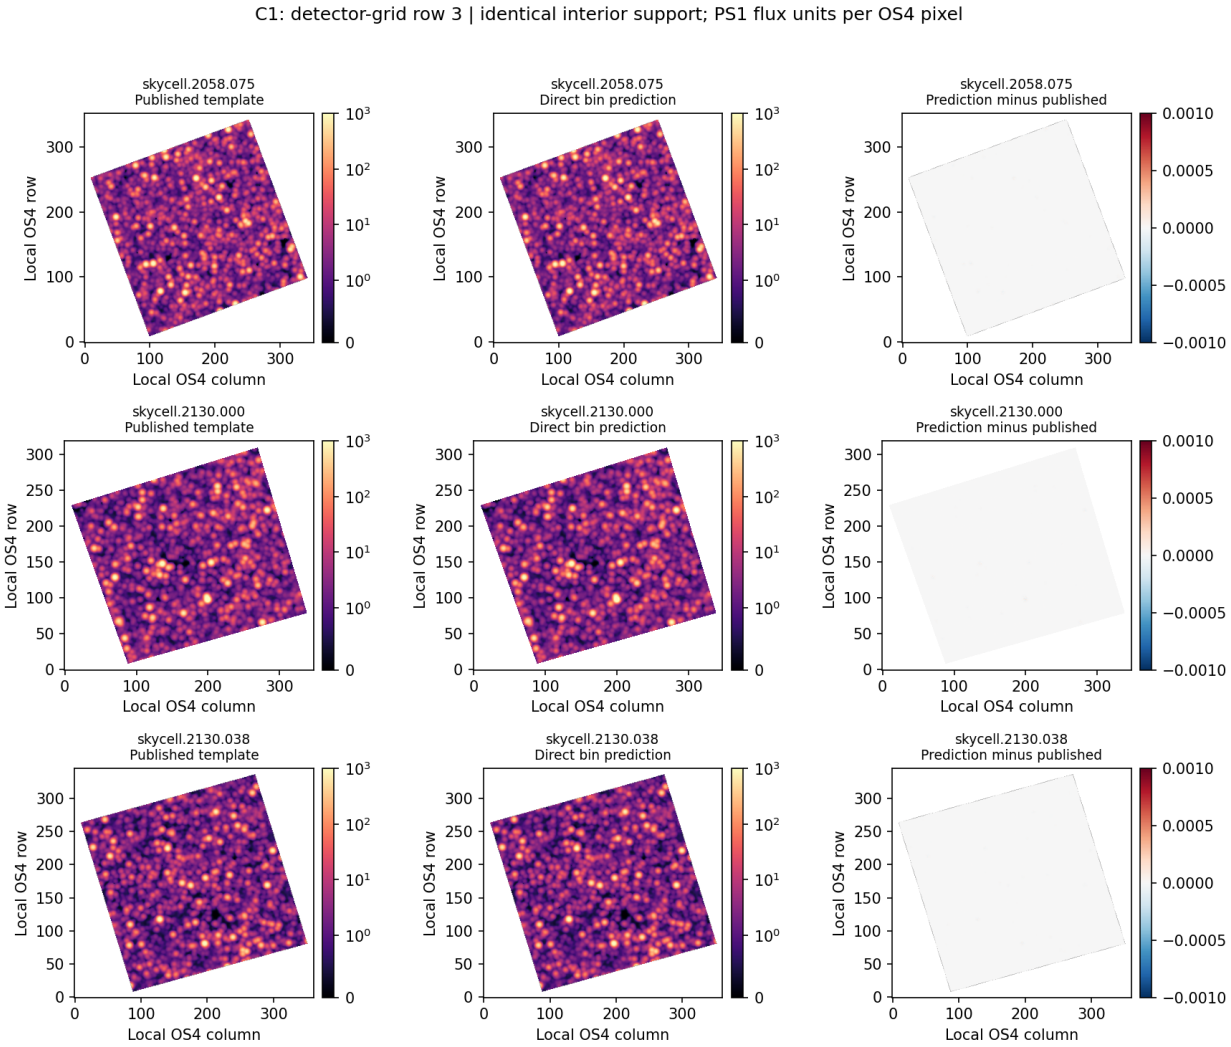


Figure 37. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2058.075	67,260	1.96×10^{-6}	0
skycell.2130.000	63,289	1.94×10^{-6}	0
skycell.2130.038	72,889	1.46×10^{-6}	0

C4: all four original PS1 bands and source removal

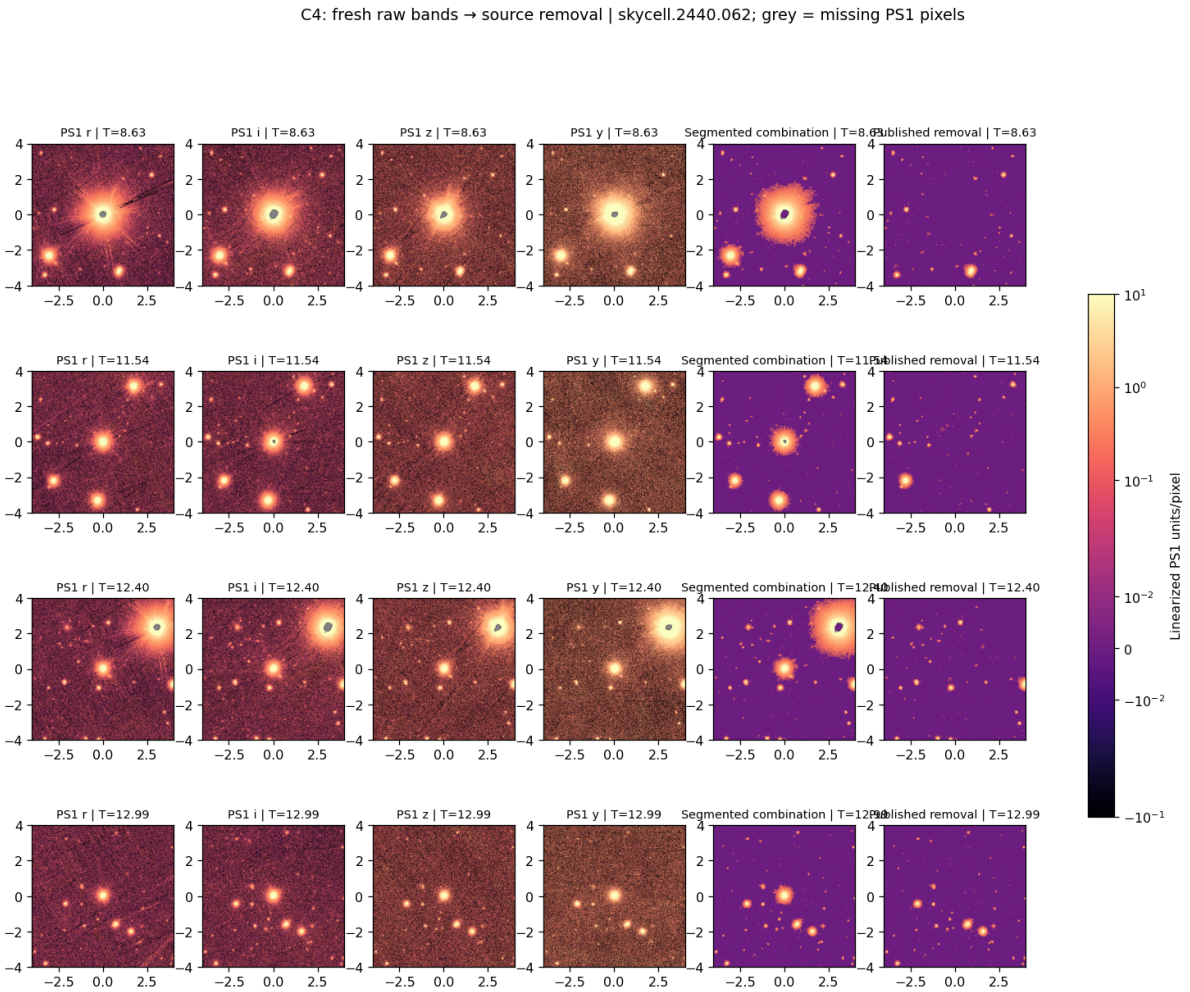


Figure 38. Four magnitude-spread examples, identical source positions. Columns show the original r, i, z, y bands, segmented combination and published removal. Grey marks unavailable original PS1 pixels; intensities in different bands differ physically. The centre-cell full-array replay is exact. The raw PS1 values have their own flux convention, not TESS electron/s.

C4: absolute stage images and paired differences

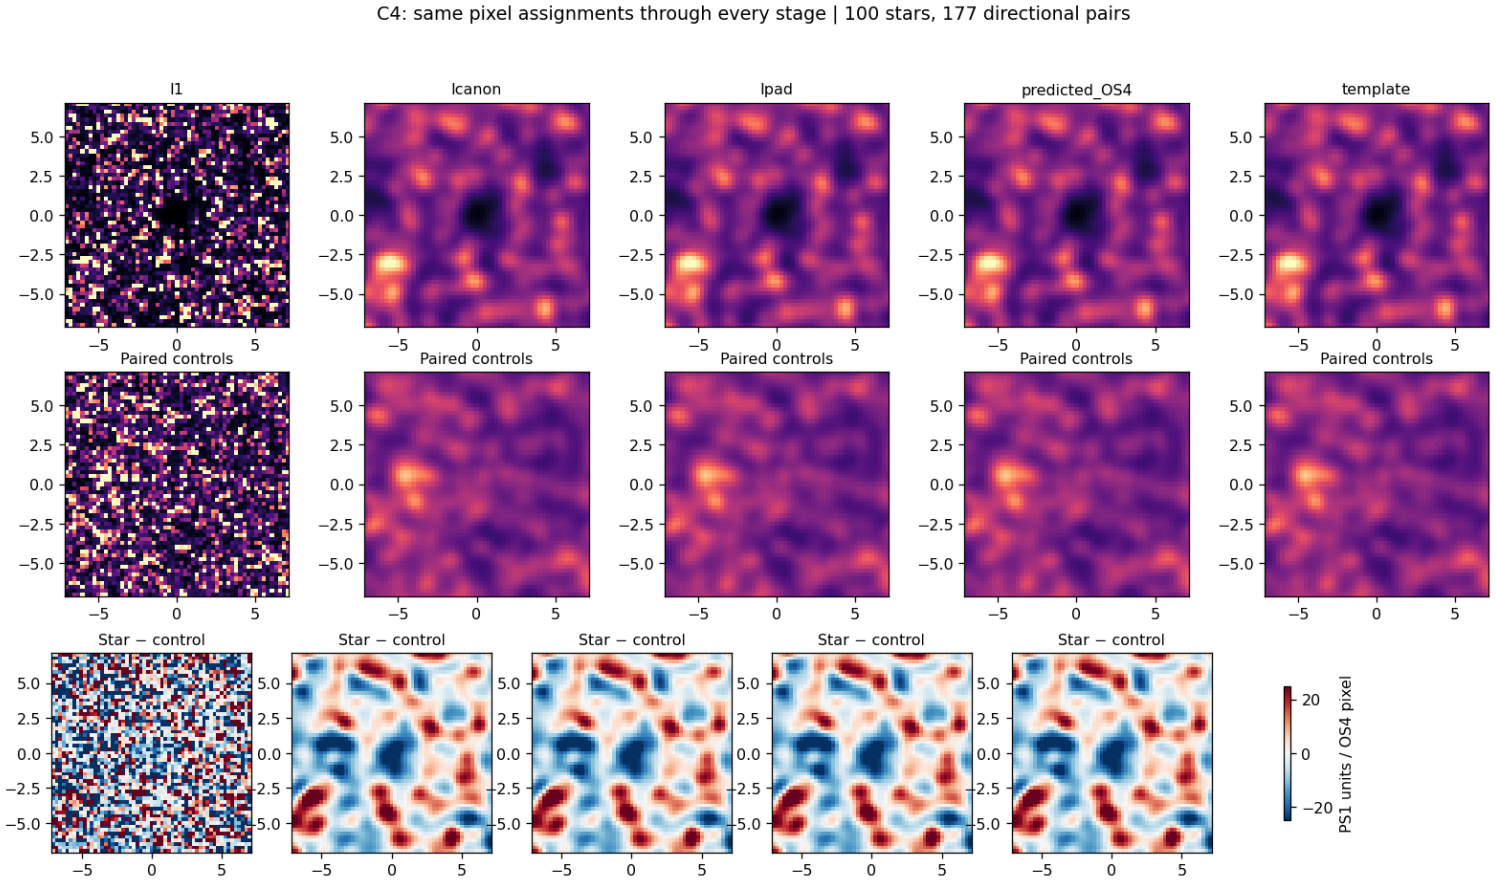


Figure 39. Top: mean at removed-star positions. Middle: paired directional controls. Bottom: their difference. Columns progress from published source-removed PS1 pixels through canonical and padding-aware convolution to predicted and published OS4 template. The central depression is already present in the first column; the later calculations propagate it. Native PS1 sums have been put on the same frozen OS4 assignments; units are PS1 flux per OS4 pixel.

C4: an individual patch and the mean saved-kernel replay

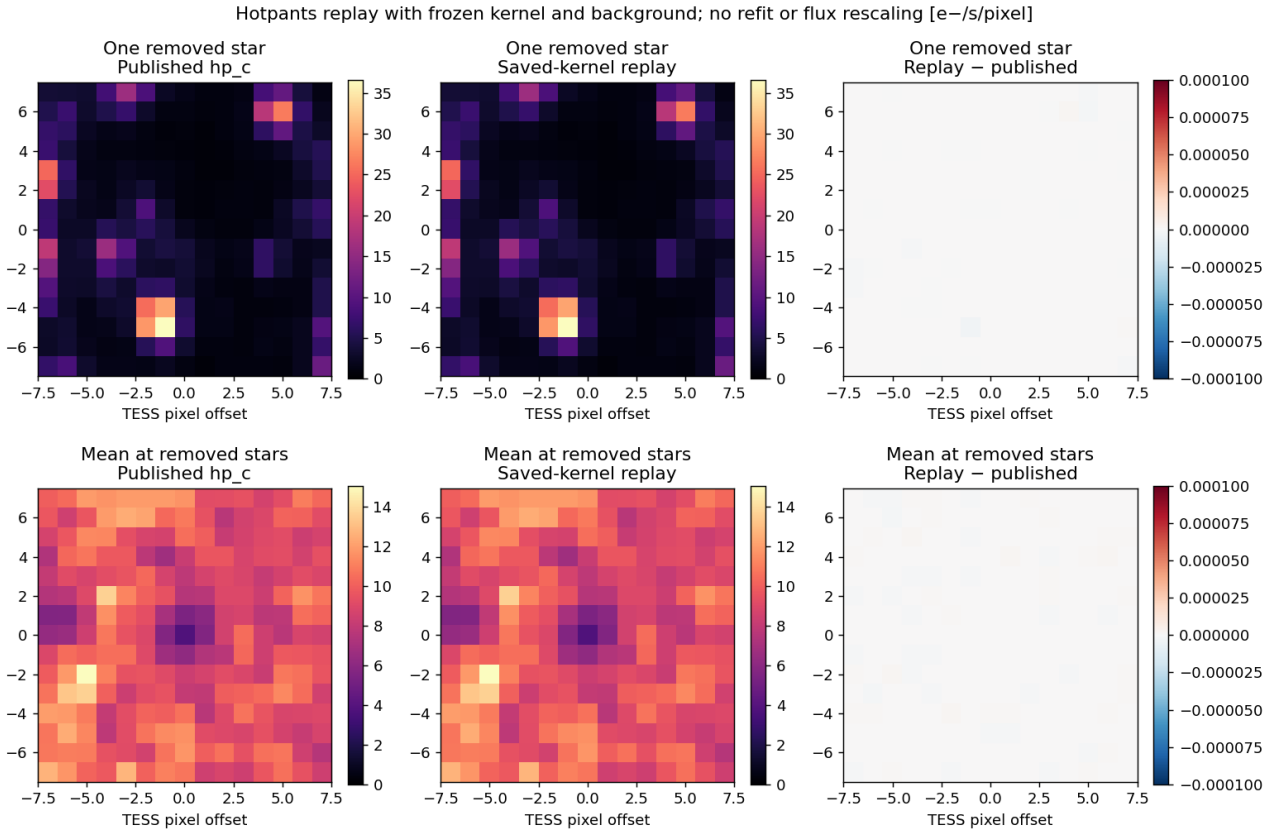


Figure 40. Published final template, replay with the frozen field kernel/background, and their difference. The first row is an individual removed-star patch; the second is the mean over tested star positions. The residual has its own narrow scale. No fit or flux rescaling was used; units are electron/s per native pixel.

C4: untouched-sky selection null

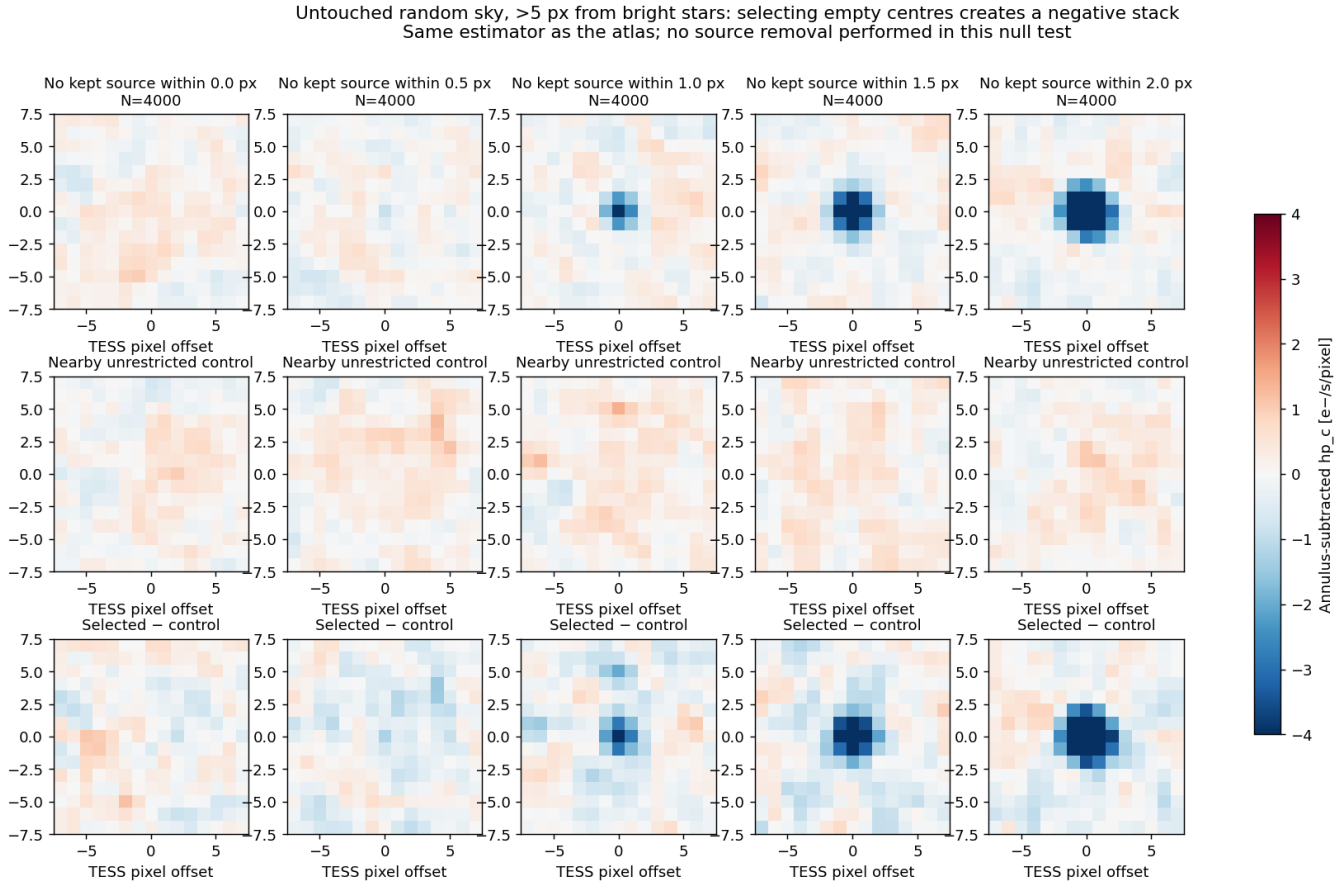


Figure 41. Columns change only the minimum distance to retained Gaia sources. Rows show selected unchanged locations, their nearby unrestricted controls and the difference. Both centres are away from bright stars. All images use the same scale in electron/s/native pixel after local annulus subtraction. A depression can be generated by changing the source population sampled by the estimator. Strong off-centre structures in the F2 controls account for its larger uncertainties.

C4: real stars, matched controls and both contrasts

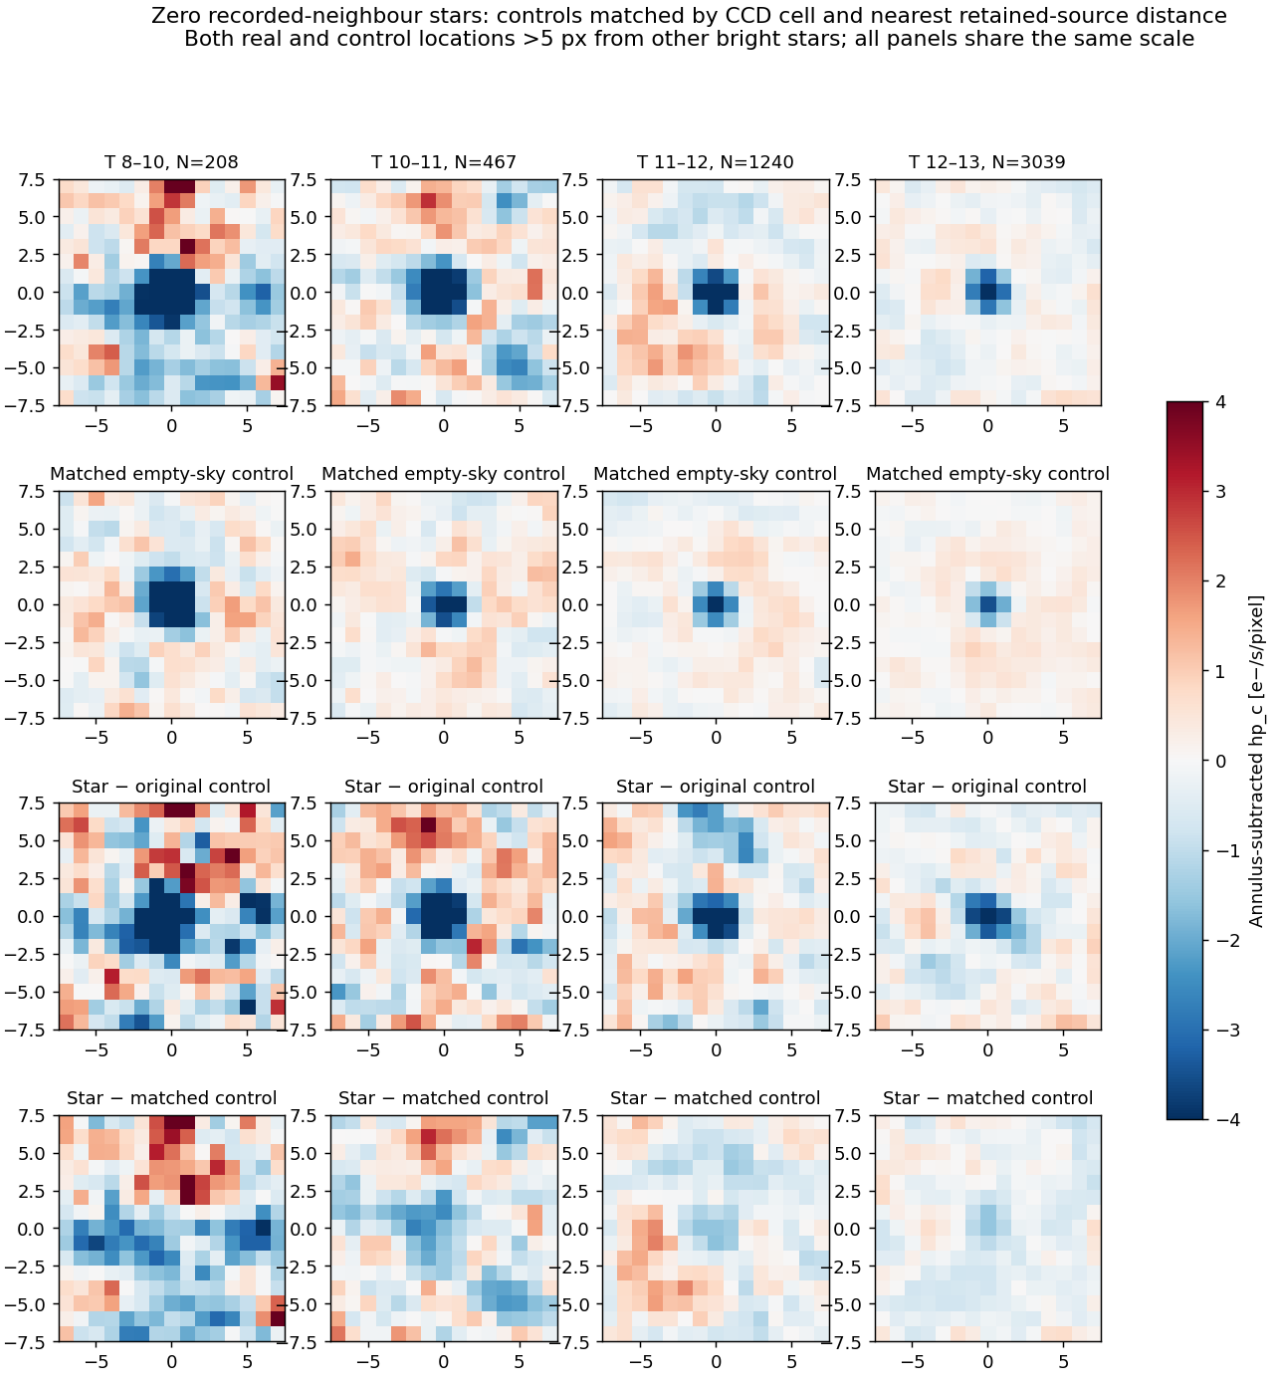


Figure 42. Columns are the four catalogue-magnitude bins. Rows show absolute annulus-subtracted star stamps, matched controls, the original contrast and the matched contrast. Matching never uses central template intensity. All panels share one scale, in electron/s/native pixel. The complete table in the main text reports aperture sums of these same difference images.

C4: detector-grid row 1, direct template reproduction

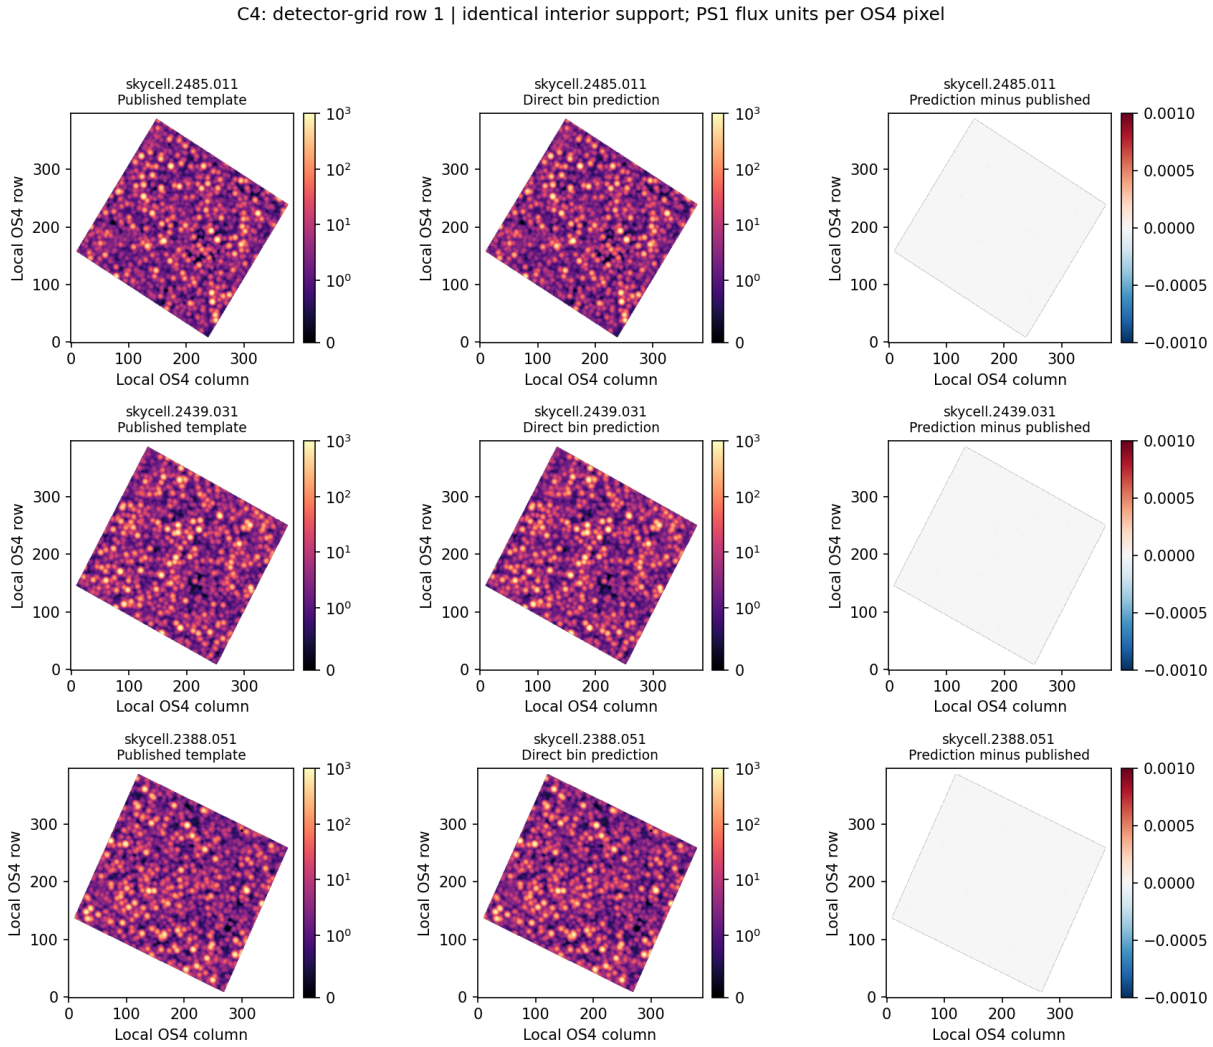


Figure 43. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2485.011	73,751	2.14×10^{-6}	0
skycell.2439.031	75,723	2.39×10^{-6}	0
skycell.2388.051	79,304	2.14×10^{-6}	0

C4: detector-grid row 2, direct template reproduction

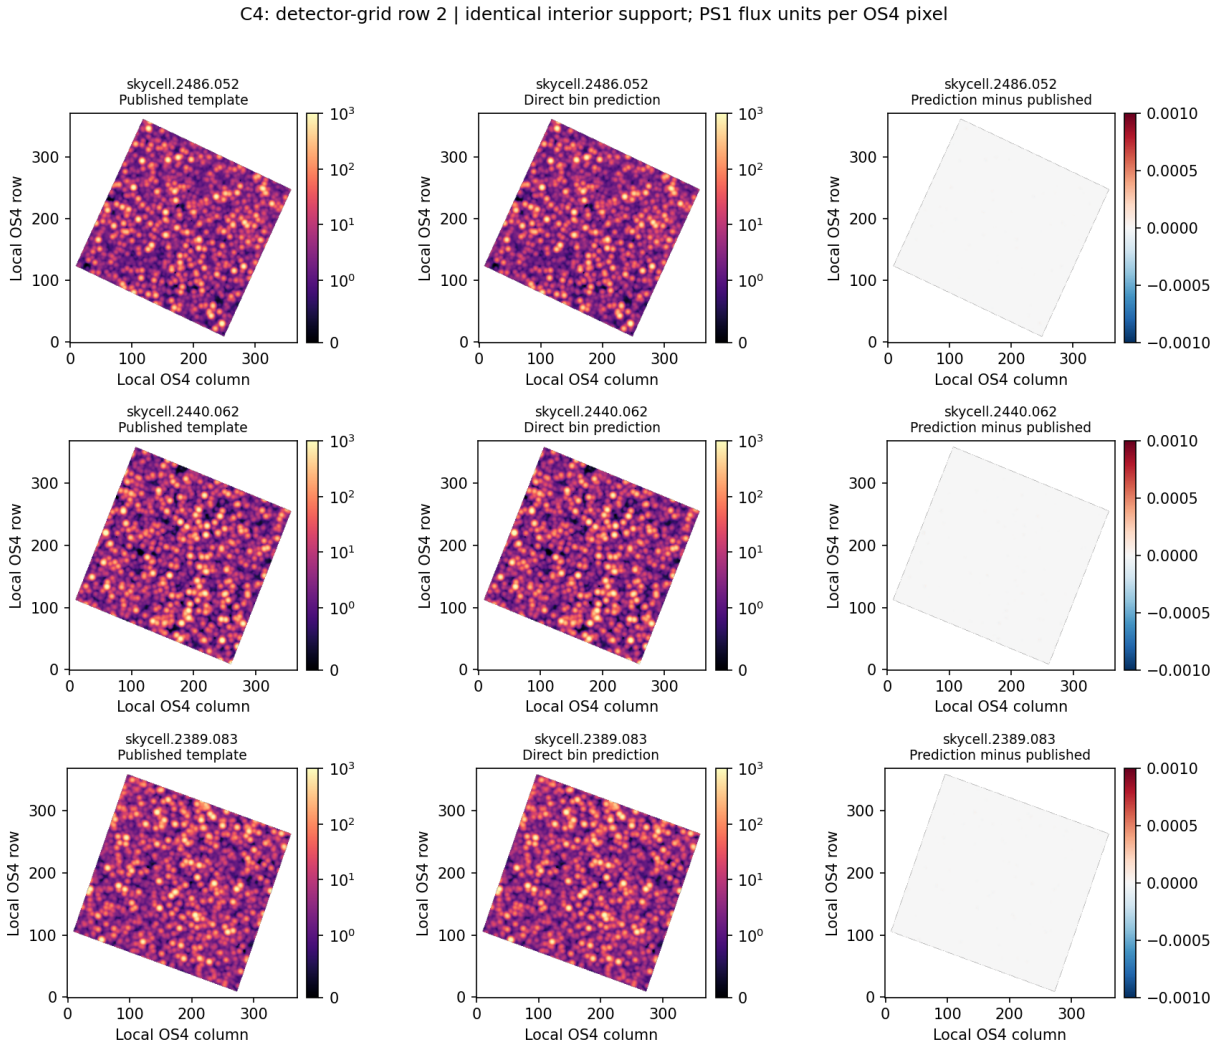


Figure 44. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2486.052	69,635	2.03×10^{-6}	0
skycell.2440.062	71,835	2.65×10^{-6}	0
skycell.2389.083	75,187	2.27×10^{-6}	0

C4: detector-grid row 3, direct template reproduction

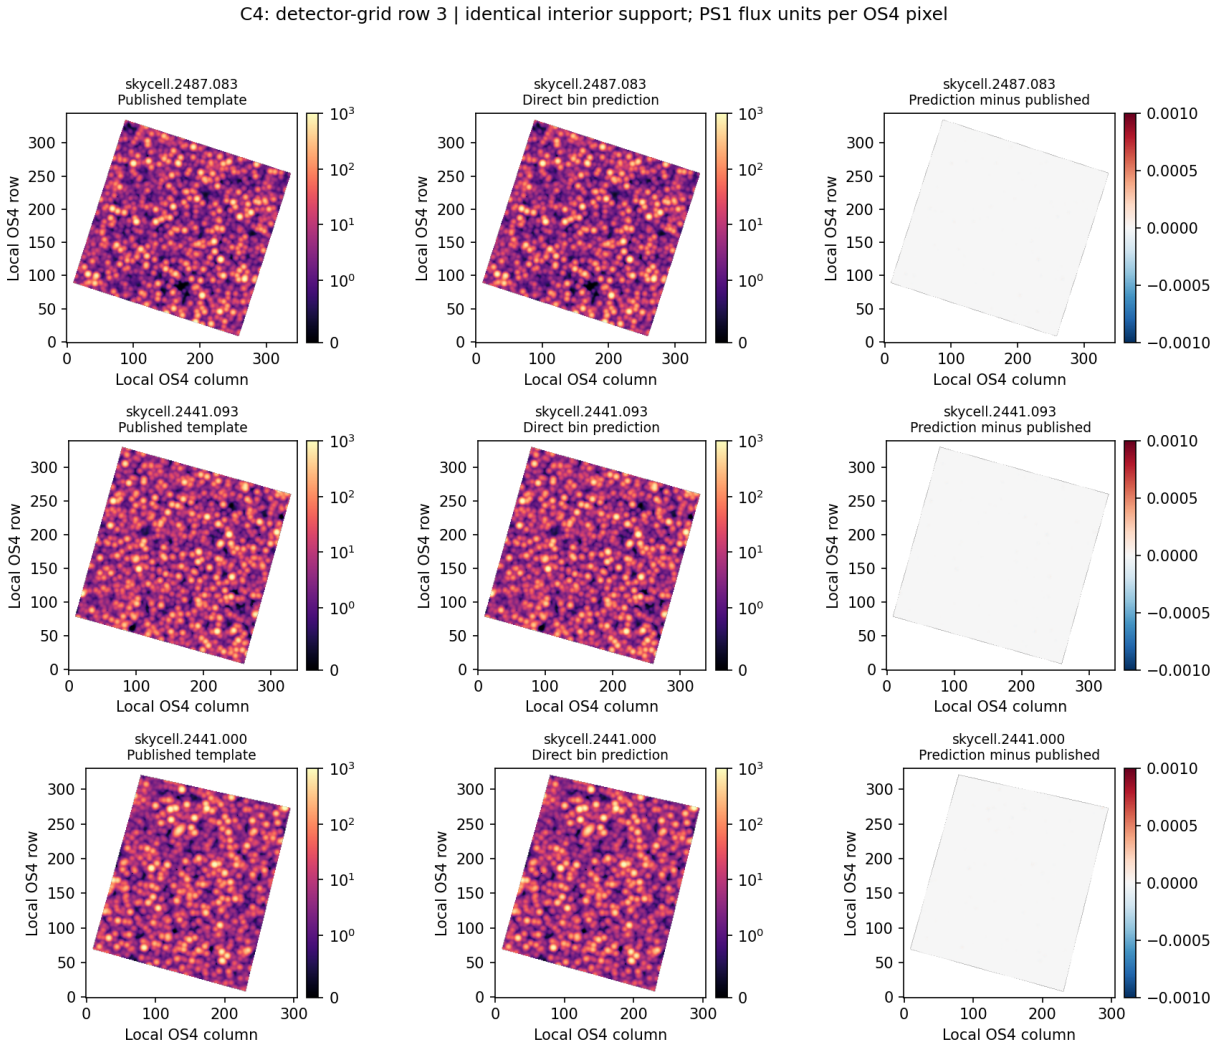


Figure 45. Left: published OS4 template. Middle: direct binning of the actual padding-aware PS1 inputs with bit 12 suppressed. Right: prediction minus published template. The same valid interior support is used in each row. Absolute maps share a PS1-unit scale; differences share a narrow symmetric scale of $\pm 10^{-3}$ PS1 units per OS4 pixel. The origin of the displayed rectangular patch is local to its sampled skycell.

Skycell	Interior pixels	Binning RMS	Assembly max
skycell.2487.083	67,479	2.32×10^{-6}	0
skycell.2441.093	67,905	2.21×10^{-6}	0
skycell.2441.000	60,081	2.56×10^{-6}	0

B Reproducibility and records

The final numeric comparison is read from `multifield/summary.json`; tables in this document are generated from that file and individual cell QA records, rather than transcribed from chat. The frozen diagnostic revision is `916d59a` on the nested dev-repository branch `codex/ps1-stage-trace`. The report-generation source is committed separately with the report source. All large arrays and figures remain on `/astro`.

- Run root: `/astro/armin/koji/syndiff/dev_runs/epsf_stage_trace_20261002/`.
- Five-field manifests, logs and numeric records: `/astro/armin/koji/syndiff/dev_runs/epsf_stage_trace_20261002/multifield/`.
- Per field: `inputs.json` (immutable input hashes), `DONE.json`, `tables/cell_inventory.json`, `tables/raw_replay.json`, `tables/stage_summary.json`, `tables/hpc_replay.json`, `tables/matched_controls_brightmatched.json`, and the underlying images/arrays.
- Earlier investigation and corrections: `/astro/armin/koji/syndiff/dev_runs/epsf_atlas_20261001/`.
- Diagnostic worktree: `/home/kshukawa/syndiff_pipeline/dev/.claude/worktrees/codex-ps1-stage-trace/`.
- Pinned processing source: `/home/kshukawa/syndiff_pipeline/.claude/worktrees/paper-pin/`, full revision `648c0bcf4b559691177f7b4764f721ac54aebdc8`.
- Local illustrated canvas: `/home/kshukawa/syndiff_pipeline/doc/ps1_stage_trace_canvas/index.html`.

The original numerical checks were run in the `syndiff` environment, with BLAS/OMP/Numba thread count 1 and two Dask workers per field. Each field used the same command with its own label:

```
mamba activate syndiff
python .../diagnostics/ps1_multifield.py --field F1
# Repeat with F2, S22, C1 and C4.
python .../diagnostics/ps1_multifield_report.py
```

The source bundle for this document includes `report.tex`, the embedded image files and a build command. Numeric source hashes are recorded in `report_manifest.json`. PDF compilation does not rerun any scientific measurement.

Conclusion, restated with its scope: direct tracing supports the implemented processing of the tested five-field interiors. The astrophysical meaning of the smaller residual contrasts remains unresolved. Paired known-flux injection is the appropriate next test.