

Image Calibration and Instrument Signal Removal for the First Year of the LSST

ALEXANDER BROUGHTON ¹, AGNÈS FERTÉ ², CHRISTOPHER Z. WATERS ³, ELI S. RYKOFF ¹,
SEAN PATRICK MACBRIDE ⁴, ANDRÉS A. PLAZAS MALAGÓN ^{2,1}, ROBERT H. LUPTON ³, AND ELANA K. URBACH ⁵

¹*Kavli Institute for Particle Astrophysics and Cosmology, SLAC National Accelerator Laboratory, 2575 Sand Hill Rd., Menlo Park, CA 94025, USA*

²*SLAC National Accelerator Laboratory, 2575 Sand Hill Rd., Menlo Park, CA 94025, USA*

³*Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA*

⁴*Physik-Institut, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland*

⁵*Department of Physics, Harvard University, 17 Oxford St., Cambridge MA 02138, USA*

(Dated: August 25, 2026)

ABSTRACT

The Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) requires calibration products that provide uniform, stable, and accurate photometric and astrometric performance across the 3.2 gigapixel focal plane and throughout the 10-year survey. This paper details the algorithms and workflows used to produce instrument calibrations and remove instrumental artifacts for Data Preview 2 (DP2)—the first end-to-end processing demonstration using on-sky data with the LSST Camera (LSSTCam). We describe the verification, acceptance, and certification framework used to assess calibration quality and quantify residual systematics. We show baseline metrics on calibrated science images to evaluate the robustness of the calibration and instrument signature removal (ISR) data processing pipelines for DP2. Finally, we summarize the known limitations observed in DP2 production and outline expected algorithmic improvements for the first public LSST data release (Data Release 1, DR1).

1. INTRODUCTION

The Vera C. Rubin Observatory will conduct the 10-year Legacy Survey of Space and Time (LSST, [Ž. Ivezić et al. 2019](#)), imaging roughly 18,000 deg² of the southern sky in six broadband, optical-near-infrared (NIR) filters (*ugrizy*). Achieving the LSST science goals requires calibrating the instrument with high precision and uniformity so that raw detector readouts can be transformed into science-ready images and catalogs. Errors or non-uniformities in image calibration can introduce survey-scale photometric and astrometric systematics that propagate into downstream measurements and may be degenerate with astrophysical or cosmological signals ([J. H. Esteves et al. 2023](#)).

Key challenges to delivering this level of calibration include:

1. **Instrument complexity:** The LSST Camera (LSSTCam) is the largest digital camera built for astronomy. Its 3.2-gigapixel focal plane consists of 205 fully depleted, backside-illuminated charge-coupled devices (CCDs) from two vendors (Teledyne e2v and Imaging Technology Laboratory, hereafter E2V and ITL). Detector-to-detector and

amplifier-level variations require calibration and processing algorithms robust to heterogeneous sensor behavior.

2. **Data volume and cadence:** LSSTCam typically acquires 30 s exposures with ≤ 3 s read-out, yielding about 1000 exposures per night and nightly data rates of approximately 10 TB. The need for fast, efficient processing speed is driven by Prompt Processing for the alert database and the active optics system.
3. **Calibration system and optics:** The calibration system (flatfield screen, LEDs, tunable laser, and the Collimated Beam Projector or CBP) must provide uniform and repeatable illumination over a large focal plane through a fast converging ($f/1.2$) optical system.
4. **Coupled instrumental effects:** Instrument signature removal (ISR) must correct artifacts in the reverse chronological order in which they imprint in the photon transfer chain. Calibration derivation follows the same ordering, since characterizing a given artifact requires that all effects im-

printed before it have already been calibrated and removed. This is further complicated by partial degeneracies among instrument artifacts, which necessitate iterative or “bootstrapped” approaches to converge on self-consistent calibration products.

5. **Evolving conditions and cadence:** Changing environmental conditions, multi-epoch calibration datasets, and evolving observing strategies require calibration production to be robust to non-uniform survey and telescope operations. The datasets discussed in this paper were collected before the dome was completed, before stabilization of thermal regulation systems, and before finalization of the precise detector operating mode.
6. **Verification, provenance, and operational reproducibility:** Calibration products must be accompanied by quantitative diagnostics for rapid data quality assessment, cross-time and cross-version comparisons, and clear acceptance and certification decisions for downstream processing. Products must be reproducible, versioned, and continuously validated as instruments and algorithms evolve.

The LSST Data Management Science Pipelines (Rubin Observatory Science Pipelines Developers et al. 2026) implement this functionality using production pipelines and supporting tools to generate, verify, and publish calibration products for routine processing. The design benefits from lessons learned in LSST Data Preview 1 (DP1) (Vera C. Rubin Observatory Team 2026) and precursor surveys, including the Hyper Suprime-Cam Subaru Strategic Program (S. Miyazaki et al. 2012; H. Aihara et al. 2022; J. Bosch et al. 2018) and the Dark Energy Survey (T. D. E. S. Collaboration 2005; T. Diehl & F. T. D. E. S. Collaboration 2012; E. Morganson et al. 2018).

The first end-to-end demonstration of Rubin data processing on a substantial, internally consistent LSSTCam dataset is *Data Preview 2* (DP2). DP2 covers about 2,500 deg² in all six filters and includes commissioning and pre-survey operations data. This paper describes the calibration-production algorithms and the operational approach used to generate DP2 calibration products.

Our goals are to (1) make the DP2 calibration products interpretable to science users, (2) quantify our confidence in the quality of released data products, (3) provide a reference for future cross-survey synergies and science analyses, (4) make the processing reproducible for pipeline developers and operators, and (5) provide a

reference point for expected improvements on the path to the first LSST public data release.

This paper is organized as follows. Section §2 introduces LSSTCam, its CCD readout architecture, and the electronic detector model that organizes calibration production and ISR. Section §3 describes the DP2 calibration datasets, the calibration system (CalSys), and the criteria used to select input exposures and define calibration epochs. Section §4 describes calibration product production and application, including the algorithms and data products used in DP2. Section §5 integrates how those products compose during instrument signature removal for DP2 science data. Section §6 discusses calibration-epoch differences and residual systematics relevant to DP2. Section §7 describes how calibration products are verified, accepted, certified, and deployed, and §8 summarizes known limitations for DP2. Section §9 summarizes the DP2 calibration quality and outlines expected improvements toward Data Release 1 (DR1).

2. THE LSST CAMERA (LSSTCAM)

The Vera C. Rubin Observatory’s camera LSSTCam SLAC National Accelerator Laboratory & NSF-DOE Vera C. Rubin Observatory (2025) is the largest digital camera ever built to date. Its design, and the technical progress it demanded, were developed in order to meet the requirements set by the ambitious goals of LSST, as described in A. Roodman et al. (2024). We have thus created an ISR approach specifically for the analysis of LSSTCam images. Figure 1 shows an example of the results from our ISR pipeline on an LSSTCam exposure, showing the corrections of large effects (such as the amplifier offset or crosstalk) as well as more subtle effects (such as the brighter fatter effect) described in details in following sections. We summarize here key points about LSSTCam and properties of its CCDs relevant to LSSTCam calibration in Rubin Data Management (DM).

2.1. LSSTCam CCDs

LSSTCam is composed of 189 science CCDs organized in rafts of 9 CCDs (3x3), in a focal plane mosaic presented in CTN-001 add citation. Each raft is composed of 3 Readout Electronics Boards (REB) providing the electronics of three neighboring CCDs. In addition, 2 star guider CCDs and one split-wavefront CCD are at each four corners of the focal plane mosaic. We will focus only on the science CCDs in the following as no requirements are set on the calibration of the corner sensors within DM, although some calibrations are provided for the star guiders.

LSSTCam CCDs are made of 16 segments each, 8 at the top and 8 at the bottom of the CCD, read out




Figure 1. The left panel shows a raw image of detector 81 (e2v type) of exposure 2026010700083, and the right panel shows the resulting image from processing it through the DP2 ISR pipeline. Some corrections operated by ISR are visually noticeable such as a saturation and edge bleed masking (green mask), detector assembly, and crosstalk masking (red mask on the right).

through an amplifier. Segments are commonly referred to as amplifiers and will be in the rest of this paper. Each amplifier is 2048 pixels high and 512 pixels wide for e2v and 2048 pixels high and 576 pixels wide for ITL. Figure 2 shows a diagram of the amplifier geometry and direction of electrons transportation in LSSTCam CCDs. The overscan regions are virtual pixels read to estimate charge transfer inefficiency (CTI), see 4.6 for my details. Figure 2 show an example of the signal read in the overscan region for ITL and e2v amplifiers.

2.2. Electronic Detector Model for LSSTCam

The calibration and instrument signature removal (ISR) pipeline removes instrumental effects from science data. For each effect observed in LSSTCam, the Rubin Science Pipelines define a calibration-generation task that produces the corresponding data product, and a correction task that applies that product during ISR. We have built on the ISR approach of past optical surveys [here add citations for ISR procedures for other CCD optical experiments](#) to develop an ISR pipeline specifically designed for LSSTCam. We have defined the electronic model of LSSTCam and associated effects in the photon to ADU signal chain. Figure 3 depicts the different steps in this chain and their physical ordering from the detector to the Readout Board (REB). [here add any citations that supports this description](#).

Calibration and ISR tasks are organized as a pipeline based on our current model of the detector physics. The calibration production tasks are run in chronological order with respect to the photon transfer and data-acquisition chain, and ISR correction tasks are applied

in the opposite order to the sequence in which effects accumulate along the signal path.

The algorithms and processing steps required for LSSTCam calibration and ISR were developed from measurements and analysis during electro-optical testing on the laboratory test stand at SLAC National Accelerator Laboratory, at the summit, and on the telescope in-dome during commissioning. [A. Roodman et al. \(2024\)](#) gives a detailed account of system characterization work and results from the laboratory campaigns. Figure 3 summarizes the instrumental effects identified in the photon transfer chain and their order.

3. DATA

Summarize the DP2 context relevant to calibrations: instrument, filters, observing period/epochs, detector layout, available exposure types (bias/dark/flat/science), and definitions of calibration epochs. Optionally add science context.

3.1. LSSTCam and CCD status for DP2

3.2. Calibration System (CalSys)

LSST uses four calibration systems to construct calibrations. Three of which (the flat-field illuminator, the collimated beam projector, and the Auxiliary telescope) will be used continuously throughout the LSST, while one calibration system (the Camera Calibration Optical Bench) was used to characterize and qualify the LSSTCam CCDs before integration to the Simonyi Survey Telescope at Rubin Observatory. [add citation to Seans SPIE paper](#).




Figure 2. Here is an exposure in e2v amplifier and in ITL amplifier with rectangles around different regions (overscans, serial register) and graphs of the signal in their overscan regions.

3.2.1. Camera Calibration Optical Bench

To create optical scenes for characterization, the Camera is placed in-line with an optical scene projector called the Camera Calibration Optical Bench (CCOB) (Y. Utsunomiya et al. 2024). The CCOB-wide beam is designed to project a uniform monochromatic light source onto the focal plane, thanks to one of six LEDs, each corresponding to a central wavelength in the LSST filter bandpass. In-lab tests with the CCOB are described in [here add citation](#), the so-called 'run 7' was the final campaign of testing of LSSTCam at SLAC, calibrations produced during that run were used as the initial LSSTCam on-sky observations.

I am not sure how much we want to talk about the CCOB here. If there is a specific focus (qualifying the CCDs before integration), I can revisit this

3.2.2. Flat Field Illuminator

The Vera C. Rubin Observatory's Flatfield Calibration System is designed to provide a stable, uniform, and well-characterized illumination required to measure the absolute and relative instrumental responses of the LSST Camera. Operating in tandem with the CBP and AuxTel, this system aims to satisfy a photometric precision of 5 millimag in repeatability and 10 mmag in spatial uniformity across the sky. **Cite a flatfield calibration system paper** The physical architecture of the flatfield system is driven by the fast ($f/1.2$) optical configuration of the Simonyi Survey telescope, which demands exceptional control over both the angular and spatial

uniformity of the calibration wavefront. The system directs light toward a large 10.27-meter diameter Lambertian calibration screen that is tilted to 23 degrees to align precisely with the optical axis of the telescope. Because the calibration hardware must co-rotate dynamically with the observatory dome during operations, the primary illumination sources, environmental enclosures, and projection optics are mounted directly onto the dome structures rather than a static mount.

To deliver both broadband and monochromatic illumination during the LSST, the system integrates two independent illumination subsystems within a single centralized optical projector, positioned at the geometric center of the calibration screen structure. Broadband "whitelight" illumination is provided by a specialized module featuring a collection of 10 high-power LEDs selected to span the complete *ugrizy* filter passbands, with one LED within the *u* passband with the other five passbands each getting coverage from two LEDs. For monochromatic calibration, the system incorporates a tunable laser system that provides continuous wavelength coverage from 300 to 2600 nm at a 1 kHz pulse rate. This laser can deliver narrow tuning steps between 0.05 and 1 nm and a spectral linewidth under 1 nm FWHM. **Cite a paper by Parker on this**

The data products generated by these dual operating modes support separate but complementary calibration goals within the Rubin DM pipeline, to ensure uniformity across the LSSTCam focal plane. In the DP2 dataset, the flatfields used during ISR are constructed with a weighted combination of the two LEDs




Figure 3. Photon transfer and data-acquisition model for the LSSTCam detectors and electronic readout board (REB). Instrumental effects seen in images are shown as dashed boxes; arrows indicate where each effect enters the chain. Low-signal nonlinearities in LSSTCam’s detector response have been observed, but their origin is not yet established; that contribution appears as a disconnected box.

within the passband (or the single LED in the case of u band). These data products enable the correction of detector-level systematics, including gain variations, charge transfer inefficiencies, pixel-level defects, and sensor distortions induced by the brighter-fatter effect. The monochromatic flatfield data products from the tunable laser are not used in DP2, but will be utilized in future data releases to map the system’s chromatic throughput as a function of focal plane position.

3.2.3. Collimated Beam Projector

The Collimated Beam Projector (CBP) is a dedicated, in-dome photometric calibration instrument to measure the system throughput of focused light to millimag precision **Cite a CBP paper**. The CBP consists of a modified Schmidt reflector telescope fed by the tunable laser to produce a collimated beam of known flux. By co-moving the CBP and Simonyi telescope, we are able to place this beam anywhere on the primary mirror and have a focused spot anywhere on the focal plane. The CBP can switch between five different masks at the focal plane, allowing for different spot patterns to be projected onto the focal plane. For DP2, a mask that projects once spot per CCD was used to construct the

crosstalk correction. However, the CBP is also capable of measuring QE variations across detectors and filter transmission curves, and will be used for these purposes in DR1. **cite CBP SPIE paper Do we want an image here?**

3.2.4. The Auxiliary Telescope

The Vera C. Rubin Observatory’s Auxiliary Telescope (AuxTel) is an independent 1.2-meter telescope located on the same ridge as the Simonyi telescope. It has a slitless spectrograph that measures atmospheric transmission simultaneously with the LSST survey observations. It was in operation during much of the DP2 period, although the data taken was not incorporated into the DP2 photometric calibration. Its measured atmospheric parameters are expected to be used in future data releases.**add AuxTel citations here.**

any other instrumentation used for input datasets.

3.2.5. Limitations to Calibrations during LSST Y1

Describe any system-level constraints that influenced calibration strategy.

Describe the LSST Camera and any technical limitations (bad detectors/amps, CCD power cycling and its affect on existing calibs).

3.3. *Data selection and corresponding calibration epoch*

Describe the data selection process for the DP2 dataset and calibration epoch. See 4.

4. CALIBRATION PRODUCTION AND APPLICATION

Each subsection will describe the calibration product, its construction algorithm, and the corresponding ISR correction task. We may also include a table of the data product and the name of the associated calibration task. Figure 5 summarizes calibration-product dependencies. The integrative ISR pipeline for DP2 science data, including the order in which corrections are applied (Figure 14), is described in §5.

4.1. *Curated Calibrations*

4.2. *Auxiliary Data Products*

4.3. *Crosstalk*

As described above (XX), large astronomical sensors are read out through multiple amplifiers in parallel, reducing the time needed to completely read the sensor. One downside of this design strategy is the possibility of crosstalk between the different amplifier channels. In LSSTCam, this crosstalk occurs due to electronic coupling between the channels in the cable connecting the detectors to the REB. The result of this coupling is that the data for a particular amplifier shows a scaled copy of neighboring amplifiers added to the image. As all amplifiers are clocked through the readout process simultaneously, this scaled copy is aligned with the location of the readout corners. This may make the crosstalk signal visible in the target amplifier flipped relative to the source amplifier.

With a known set of crosstalk coefficients between a given source and target amplifier, then correcting for this additional signal is easy: each source amplifier that has a non-zero coefficient is flipped to match the target amplifier, scaled by that coefficient, and added to the results from all other source amplifier to create a subtrahend image that contains all of the crosstalk signal observed on the target amplifier. By repeating this for all amplifiers, we can build up a full-detector subtrahend image that can then be subtracted. As this correction may be imperfect, we set a CROSSTALK mask plane for all locations where the subtrahend image is greater than 1 electron. This provies a way for the downstream object detection and measurement code to identify sources

that may have been significantly altered by the crosstalk correction.

As this crosstalk occurs in the cable connecting the detector to the REB, prior to being digitized into ADU, the images need to be gain corrected prior to the crosstalk correction to ensure they are in electron units. This naturally places the crosstalk correction early in the IsrTask processing, which is important for an additional reason. Very bright sources can create bleed trails that deposit charge that is not fully clocked out, and that slowly leaks out as subsequent rows are read out. These bleed trails can extend into the parallel overscan region, which can in turn create crosstalk signal in neighboring amplifier. Performing crosstalk correction prior to the parallel overscan correction prevents these amplifiers from oversubtracting the affected columns.

The crosstalk coefficients used during commissioning were generated in the lab at SLAC prior to the camera being shipped to the summit and were quadratic with the source flux, as discussed in A. Snyder et al. (2021) (CZW: do we need to introduce the measurement process first?). These coefficients worked well, with the median values from the ITL detectors used to correct LSSTComCam while it was on sky. However, the switch from 2s readouts to 3s readouts via a camera sequencer change invalidated these coefficients, as the change in clocking also changed the scale of the crosstalk signal, requiring this calibration to be remeasured.

This measurement used three nights (CZW: check numbers in this paragraph) of spot data taken with the CBP system. Each exposure placed one spot on each detector, with ten linearly spaced flux levels (CZW: add figure or describe levels) per pointing. After the flux sequence was completed, the CBP spots where shited to land on a different amplifier, with the process repeated in order to provide each amplifier a bright source to use to look for crosstalk signals. A total of XX exposures were used to generate these updated crosstalk coefficients, with the number of pixels used as “source pixels” (CZW: which you haven’t explained yet) shown in Figure 6.

As discussed above (CZW: ???), each of these exposures was minimally processed to correct the serial overscan, apply the gains to convert the units to electrons, and to trim and assemble the data segments. All pixels above a threshold of 30,000 electrons were then identified and used to construct crosstalk ratios $r_{ij} = T(\hat{x}, \hat{y})_j / S(x, y)_i$, where $S(x, y)_i$ is the flux value in a pixel at (x, y) in amplifier i , and $T(\hat{x}, \hat{y})_j$ is the flux value at the same “readout coordinate” $(\hat{x}, \hat{y})$ in amplifier $j \neq i$. These ratios therefore show the measured response to the source signal S at the expected lo-




Figure 4.

cation of a crosstalk source at a target location T . These fluxes and ratios are accumulated for the entire set of exposures and detectors. Prior to any further study, the ratios are clipped to ensure contaminating features, such as stray light and reflections of other spots, are removed. Any pixel with a measured ratio larger than 1.0 are excluded, as these would represent an unphysical situation in which the crosstalk signal was brighter than the source signal. The remaining signals were XX-sigma clipped based on the inter-quartile distance, under the assumption that the ratio values are normally distributed.

If the crosstalk response was perfectly linear, the average of these clipped ratios would directly yield the crosstalk coefficient between the source and target amplifiers. However, as shown in the lab study of the crosstalk, and demonstrated in Figure 7, such a simple model fails to fully remove the crosstalk, as the source brightness changes over the profile of the spot.

To provide the required crosstalk coefficients that are non-linear in source flux, a linear fit was made to the measured ratios as a function of the source flux. The intercept of this fit then provides the term that is linear in the source flux, and the slope of the fit providing a term quadratic in the source flux. CZW: is this clear?

4.4. Defects

We build a DEFECTS calibration product based on defects manually defined and automatically detected. They correspond to pixels which output is not well corrected by our ISR approach. These Defect pixels are masked and marked as BAD for the downstream analysis and interpolated over for visual inspection.

The DEFECTS calibration product is constructed by measuring defects on bias, dark, flats and then merged together. The defects detection consist of the following steps:




Figure 5. Calibration production dependency graph. Arrows indicate data or dependency flow (“ $x \rightarrow y$ ” means “ y depends on x ”). Solid outlined boxes represent datasets that are deployed/released with data products while dashed boxes represent internal or intermediate datasets that are not formally released to users.

- correct for the gradient across the focal plane, by applying a gradient correction to each amplifiers from the gradient measured in a binned image.
- a background estimate is done and removed,
- hot pixels as defined as pixels with more than 5 e-/s (scaled by the exposure time) of dark current, above 1000 ADU in bias are measured,
- cold pixels as defined as pixels with photoresponse less than 80% of the flat mean are measured,
- pixels along the midline break (the interface between the top and bottom amplifiers) in E2V detectors are set to the median of the flats (so they are not found by the defect code),
- bad columns masking.

Afterwards, these measured defects are merged together, and the edges of the detectors are masked, to




Figure 6. Number of source pixels used to measure the crosstalk coefficients for each amplifier. Note that this number drops toward the edge of the focal plane, due to difficulties properly illuminating CBP spots in these locations.




Figure 7. An example crosstalk residual on LSSTCam detector XX after correction with a linear crosstalk coefficient. The oversubtraction in the core and undersubtraction in the wings of the source illustrates that the strength of the crosstalk signal is dependent on the flux of the source pixel. The right panel shows the same spot, as corrected with the final non-linear crosstalk coefficients.

mask the so-called picture frame effect corresponding to the presence of large electric field at the edges of the CCDs.

For DP2, these defects were produced by using 'bootstrap' version of the bias, dark, flat as these calibration products need the defects themselves to be defined. This will probably be updated in future iteration of calibration products.

Additionally manual defects are large features present in LSSTCam detectors, such as the large defect in detector 0 which induces large bleeds. LSSTCam manual defects are defined in https://github.com/lsst/obs_lsst_data/tree/main/lsstCam/manual_defects.

Further defect statistical analyses, description of the defect monitoring, masking optimization are postponed to a future paper. Additional masking is applied to an exposure in the ISR processing, as described in Sec ??.

4.5. Non-Linearity

Non-linearity (NL) describes a deviation from simple proportionality when mapping amplified voltage integrated during readout to the number of analog-to-digital units (ADU) or counts. Non-linearity is signal-dependent and is often not a simple function of signal level; it can change suddenly above a large signal level that is characteristic to each readout amplifier; and it can drift with sensor temperature and with time, as slow electronic shifts in the bias level accumulate over a long calibration campaign. In practice it introduces pixel-level measurement bias (up to x% from the true signal) and invalidates Poisson counting statistics and can interfere with measurement of gain (which shouldn't change or drift with signal level). Correcting NL to better than y% is therefore a prerequisite for reliable downstream calibration products and science analyses.

We can measure non-linearity by measuring the response of pixel values to a change in illumination level. The in-dome calibration illumination system can deliver an approximately uniform flux of light across the LSST-Cam focal plane. Given a known change in some input illumination level, we should be able predict a change in the number of counts in an image, which we can then compare to the measured change counts.

We do not have independent measurements of the photodiode current produced by our calibration illumination system that we can integrate of the exposure time and use as a baseline truth of the amount of incident light while taking an exposure. The exposure time, which is precisely measured, makes a natural proxy for missing photodiode measurements. We therefore calibrate NL by fitting a relation between exposure time (Δt , abscissa) and the mean measured flat-field signal (μ , ordinate).

We measured 543×2 flat field exposures at evenly spaced illumination levels. We went as low as




Figure 8. Mean flat-field signal as a function of exposure time for a representative amplifier. Discontinuities at LED current-setting boundaries partition the data into distinct illumination groups.

100 analog-to-digital units (ADU) and continued beyond pixel saturation. At each illumination level, a pair of flats was acquired by varying both photodiode current and exposure time to achieve the targeted levels, measuring the mean signal level as $\mu_i = \text{avg}(F_i^{(1)}, F_i^{(2)})$ (in ADU). These data were acquired with illumination levels decided in random order to decouple signal-dependent non-linearity and time-dependent drifts in non-linearity.

Figure 8 shows the resulting relation for a representative amplifier. Discontinuities appear at the boundaries between LED current settings, producing different signal levels for the same exposure time. These data are partitioned into groups associated with distinct LED illumination settings. Grouping is determined by computationally searching for jumps in the ratio of mean detector signal to exposure time or via exposure meta-data. Within each group we allow a distinct proportionality constant k_j to capture the size of the jumps.

To first order, the trend between exposure time and the mean measured flat-field signal over properly scaled groups is linear and related to the intrinsic gain in each readout channel. We can fit a simple slope and intercept to this trend to the low signal points; the highest signal level in each channel above which this approximation breaks is the “linearity turnoff” (as indicated in Figure 9). Points above the turnoff are excluded from the NL fit because there is a maximum signal above which the mea-




Figure 9. Linearity turnoff for a representative amplifier. The turnoff marks the highest signal level below which the exposure-time-to-signal relation remains approximately linear.

sured counts can no longer be trusted due to hardware limitations at the output of the detector (Figure 3).

The NL model we fit to this relation is an extension of the formalism developed by P. Astier et al. (2019). We represent the deviation from linearity by an Akima spline $S(\mu)$ of the measured mean signal μ .

For each exposure i in group j we define the residual ratio

$$r_i = \frac{\mu_i - S(\mu_i) - O}{k_j D_i} \quad (1)$$

where D_i is the data used for the abscissa (exposure time), O is an optional constant offset that absorbs scattered light or other signal not proportional to D_i , and w_i are weights.

The fit minimizes a uniformly weighted ($w_i = 1$) least-squares objective over all accepted points:

$$\chi^2 = \sum_i w_i (r_i - 1)^2. \quad (2)$$

The fit is subject to the boundary conditions $S(0) = 0$ and $\langle S(\mu) \rangle = 0$ over the fit range. The first constraint fixes the origin such that zero illumination is always measured zero measured signal; the second ensures that the spline captures deviations from linearity independently from knowing the true amplifier gain.

In addition, any hardware sensitivities are absorbed by nuisance parameters that we parameterize as 1st order

linear corrections to μ before the spline is evaluated:

$$\mu \rightarrow \mu (1 + \alpha T_i), \quad (3)$$

$$\mu \rightarrow \mu (1 + \beta t_i), \quad (4)$$

where T_i is a readout temperature (referenced to the median over the campaign), t_i is the Modified Julian Date (MJD) of the midpoint of the shutter open time, and α and β are fitted coefficients that describe each sensitivity. For LSSTCam DP2 calibration production we fit both the temperature and temporal coefficients as well as the scattered-light offset O on the spline.

We could perform this fit independently for every amplifier on every detector. However, most of the non-linearity is common to all amplifiers due to similar readout architectures, with only small amplifier-to-amplifier variations on top that. In addition, each independent fit to every amplifier would pick up not just the intrinsic non-linearity in amplifiers, but also exposure-to-exposure spatial illumination variation that occurs during the calibration campaign. To better characterize this, it makes sense to perform a spline fit on each amplifier relative to both a single reference exposure and a reference amplifier on the same detector.

We can distinguish an *absolute* non-linearity in a single reference amplifier on a detector and the *relative* non-linearity of other amplifiers to that reference amplifier, the process for which is summarized in Algorithm 1. After choosing a reference amplifier and a reference exposure, we do (1) a relative spline fit on every other amplifier using the variation of its mean flat-field signal μ relative to the reference amplifier and to the μ in a single reference exposure, and then (2) an absolute spline fit on the reference amplifier using the usual from μ versus Δt . Every amplifier on a detector has a shared absolute spline fit based on the reference amplifier for that detector and a relative spline fit computed relative to a reference amplifier and a reference exposure ($S_{\text{rel}}(\mu) = 0$ for the reference amplifier).

Algorithm1 Double-spline linearity fit (single detector)

```

1: for each amplifier  $a$  do
2:    $k_j \leftarrow \text{FIND\_GROUPS}(\mu_a, \Delta t)$ 
3:    $\text{COMPUTE\_LINEARITY\_TURNOFF}(\mu_a, \Delta t)$ 
4: end for
5: choose reference amplifier  $\text{ref}$  with the highest linearity
   turnoff
6: choose reference exposure  $i_0$ , the flat whose mean signal
    $\mu_{\text{ref}} = 10^4$  on amplifier  $\text{ref}$  is nearest the target signal
   level
7: for each exposure  $i$  do
8:   for each amplifier  $a \neq \text{ref}$  do
9:      $D_i^{\text{rel}} \leftarrow \mu_{\text{ref},i} \times (\mu_{a,i_0} / \mu_{\text{ref},i_0})$  (Relative Abscissa)
10:     $\mu_i^{\text{rel}} \leftarrow \mu_{a,i}$  (Relative Ordinate)
11:   end for
12:    $D_i^{\text{abs}} \leftarrow \Delta t_i$  (exposure time)
13:    $\mu_i^{\text{abs}} \leftarrow \mu_{\text{ref},i}$  (measured counts)
14: end for
15:  $S_{\text{abs}}(\mu), O, \alpha, \beta \leftarrow \text{MINIMIZE}(\text{Equation 2}, D_i^{\text{abs}}, \mu_i^{\text{abs}}, k_j)$  ▷ Spline fit
16:  $S_{\text{rel}}(\mu) \leftarrow \text{MINIMIZE}(\text{Equation 2}, D_i^{\text{rel}}, \mu_i^{\text{rel}}, k_j)$ 

```

To determine the corrected mean signal level, we can interpolate the Akima splines to compute a correction: $\mu' = \mu - S_{\text{rel}}(\mu) - S_{\text{abs}}(\mu)$.

If each amplifier is properly linearized around its true gain, we should observe in every exposure that the illumination level is coherent between amplifiers and detectors. If the linearization is miscalibrated or biased (either because exposure time is not a good proxy for the true integrated photodiode or if constraining $\langle S(\mu) \rangle = 0$ is biased), the amplifiers and detectors will show signal-dependent fractional offsets in mean signal levels in each flat similar to a mis-estimation of gain.

We can check this how the linearized mean signal levels deviate from an overall linear relation over some set of detectors in a given exposure: For each exposure i , we form the linearized signal, μ'_i , on the reference amplifier on every detector and compute the fractional residual:

$$r_i = \frac{\mu'_i - m_i D_i}{y_i}, \quad (5)$$

where m_i is the slope of a linear fit to μ'_i versus D_i . We take the median $\langle r \rangle_i$ over a subset of focal-plane detectors (in the center of the focal plane) in that exposure and derive a multiplicative correction

$$N_i = \frac{1}{1 - \langle r \rangle_i}, \quad (6)$$

that we use to rescale the abscissa ($D_i \rightarrow N_i D_i$) and enforce that the linearization produces a coherent illumination across all detectors.

We can then re-compute the double-spline fit with the normalized abscissa on the reference amplifier. We can

repeat the cycle (normalize, spline, normalize, etc.) recursively several times, each time computing the normalization as

$$N_i = \frac{N_i^{\text{prev}}}{1 - \langle r \rangle_i}. \quad (7)$$

By including constraints from the whole focal plane, we get to take advantage of the added statistical information. After one iteration, this gives us a factor of N_x improvement in precision, with added value dropping rapidly after two iterations.

In ISR, given an image with pixel values measured in counts (F_{ij}), the final spline model can be interpolated, providing a continuous map from measured counts to the corrected counts for any pixel value within the range of 0 signal to the linearity turnoff, computed as: $F'_{ij} = F_{ij} - S_{\text{rel}}(F_{ij}) - S_{\text{abs}}(F_{ij})$.

4.6. Deferred Charge

Several distinct effects in the readout chain displace charge from its pixel of origin and leave correlated structure in the serial overscan: (1) a pixel-to-pixel fractional loss during charge transfer referred to as “global CTI”, (2) a local electronic-offset hysteresis in the output amplifier that mimics deferred charge in the first few overscan columns, and (3) localized capture and delayed release of charge by isolated trap sites in the serial register. Strictly, only global CTI is “charge transfer inefficiency” in the textbook sense, but because all three enter at the same point in the photon transfer chain (Figure 3) and produce overlapping signatures in the serial overscan, we model them together and produce a single calibration data product that we loosely refer to as “CTI” in Figure 5. These effects are measured, modeled, and corrected together using the extended pixel edge response (EPER) estimator developed for LSSTCam by A. Snyder et al. (2021).

Calibrating deferred charge uses the same input data as the linearity dataset (§4.5), estimating the residual charge in electrons left in the overscan region after integrating and reading out flat field images over a wide range of signal levels. The global CTI is quantified by a unitless fraction, but the trap size and the electronic drift-scale all carry signal-dependent physical units, so an imperfect conversion from counts back to physical electrons biases the eventual correction. This calculation therefore needs to happen in native electron units, so we need to remove signal-dependent non-linearities that occur when the signal is measured as counts and estimate a gain conversion from counts to electrons. For each input exposure, we apply an intermediate pass of instrument signature removal that applies defect masking and non-linearity correction, and converts pixel val-

ues to electrons, and also measures the overscan EPER statistics for each readout channel.

The natural way to convert from counts to electrons is to use the gain inferred from the PTC (§4.7), but the PTC needs to be CTI-corrected to get the most accurate inferred gains, creating what seems like a circular dependency between CTI and PTC calibration production. We break the cycle by bootstrapping. We create a temporary (intermediate) PTC that is not CTI-corrected, which we fit (see §4.7) to derive an initial rough estimate of the per-amplifier gain. That estimate is stored and applied during the ISR of the flat field images before calculating the EPER statistics. We find that the gains from a PTC without CTI correction are close enough to the final gains measured from a fully calibrated PTC (typically within 1% of the true gains). At this level, the rough gains are accurate enough for CTI calibration. The global-CTI fit is unit-free and so insensitive to any gain error, and the trap and local-offset amplitudes have only weak sensitivity to the few-percent gain bias that remains in the bootstrap estimate.

An example of the measured serial EPER statistics as a function of signal level for a single channel readout is shown in Figure 10. The global CTI produces an overall offset of this curve at all signal levels, the electronic offset produces a linear slope in the curve as a function of signal level, and the serial traps produce non-linear jumps that “turn on” above some signal level after a trap fills up and begins releasing charge with a characteristic timescale. There is a point at high signal levels at which we can no longer effectively transfer charge pixel to pixel, which we define as the “serial CTI turnoff” as indicated in the figure. We calculate this point as the highest mean signal for which the serial EPER remains consistent with a sigma-clipped linear fit to the low-signal points. For the majority of readout channels, the image region flat field level was not high enough to observe a serial CTI turnoff, which we then set as the maximum observed signal level.

The CTI solver simulates charge readout and each of the three effects at the amplifier image level for a given set of parameters. The model in the simulation is parameterized by a global serial CTI parameter b that is a fractional charge loss per pixel transfer, a local electronic offset described by drift scale A_L and decay time τ_L , and a spline-based serial charge-trap decay model whose size is allowed to depend on the signal in the last imaging column. We perform a least squares fit that finds the optimal set of parameters that can reproduce the measured overscan statistics trend as a function of signal level. Then the solver fits these parameters to the overscan image in stages. First, it varies only the




Figure 10. Serial EPER statistics as a function of flat-field signal level for detector/amplifier R22-S10/C04. Vertical markers indicate the PTC rolloff and turnoff thresholds (described in §4.7) and the serial CTI turnoff derived from the low-signal trend. The global CTI fit produces the overall offset of the measured curve at all signal levels.

local offset parameters, then the global CTI parameter, and then the trap capture curve from residuals between model and data. Exposures above configured flux limits are excluded from each stage (typically 10^4 electrons for the global CTI fit and 1.5×10^5 electrons for trap fitting).

The LSST Science Requirements Document requires that total deferred charge measured by EPER be kept under a total of 5×10^{-6} for signal levels below saturation level on each amplifier (citation needed). Within the population of amplifiers in active science detectors, only 95 amplifiers do not meet this requirement without ISR correction. Figure 11 shows the distribution of inferred global CTI across all amplifiers on the LSSTCam focal plane, and we show that 26 have a global CTI offset that is already above the maximum allowed threshold. The majority of these (19) are bad amplifiers that are in the detector corners (C00, C07, C10, C17) positioned near the serial register, and it is possible that what is inferred as global CTI could be a serial trap that turns on at very low signal levels (below the 10^4 electron level) and appears as an overall offset that is mistaken for global CTI.

When correcting deferred charge from an image during ISR, the deferred-charge correction takes in these fit parameters and applies the inverse operators from the model. The global CTI and local-offset fits, and the inverse correction operators applied during ISR, are the same for both ITL and E2V sensors, but serial traps are fitted only on ITL since it was observed that E2V sensors




Figure 11. Distribution of inferred global serial CTI across all science sensor amplifiers on the LSSTCam focal plane. *Blue*—Teledyne e2v (E2V) detectors; *Orange*—Imaging Technology Laboratory (ITL) detectors. The fit to the deferred charge model allows global CTI to vary up to 1.0×10^{-5} .

had small trap size that was difficult to measure beyond noise fluctuations and already negligible enough not to require correction. After applying a deferred charge correction, we can re-compute the EPER statistics on the corrected dataset and find that all amplifiers meet the science requirements.

While there can be deferred charge from a charge transfer in the parallel direction (shift row-to-row), in principle this is a difficult quantity to estimate accurately from parallel overscan statistics. EPER statistics computed in the parallel overscan region are noisy and reveal no statistically significant charge residuals, so we do not apply any parallel deferred charge correction to science images.

4.7. Photon Transfer Curve (PTC)

The photon transfer curve (PTC) characterizes how pixels respond to varying signal levels in a given wavelength band by measuring covariances between pixels in flat field exposures across a range of illumination levels. This dataset provides critical calibration information including amplifier-level read noise and gain for variance plane estimation, saturation thresholds for masking, and parameters needed to correct for non-linearity, charge transfer inefficiency (CTI), and the brighter-fatter effect. All of these parameters can differ at the amplifier level.

Let F_1 and F_2 denote the two flats at a given illumination level, and let $\mu = \text{avg}(F_1, F_2)$ denote the mean signal level (in ADU) over the image and away from the detector edges. The covariance between a pixel $\mathbf{x} = (0, 0)$

and a neighbor at pixel $\mathbf{x}' = (i, j)$ is:

$$C_{ij} = \frac{1}{2} \text{Cov}(F_1(\mathbf{x}) - F_2(\mathbf{x}), F_1(\mathbf{x}') - F_2(\mathbf{x}')). \quad (8)$$

Pre-model fit—Preliminarily, we fit only the variance to an approximate exponential relation derived in P. Astier et al. (2019) (their Eq. 16).

$$C_{00}^{\text{exp. approx}}(\mu) = \frac{1}{2g^2 a_{00}} [\exp(2a_{00}\mu g) - 1] + \frac{n_{00}}{g^2}. \quad (9)$$

In this expression, covariances are contributed by a gain g (electrons/ADU), noise pedestal n_{00} (electrons²), and a single brighter-fatter coefficient a_{00} (1/electrons) which has the effect of decreasing the image variance.

Near saturation, the variance deviates from this functional form, exhibiting a gradual rollover rather than a sharp cutoff due to electrostatic effects that are difficult to model analytically. On top of the exponential approximation, we introduce an empirical saturation term:

$$C_{00}^{\text{sat}}(\mu) = \begin{cases} -e^{-(\mu-\mu_0)/\tau} & \mu \geq \mu_0 \\ 0 & \mu < \mu_0 \end{cases}, \quad (10)$$

where μ_0 marks the onset of the PTC “rolloff” effects and τ controls the shape of the rollover. The total preliminary fit model is $C_{00}^{\text{prelim.}}(\mu) = C_{00}^{\text{exp. approx}}(\mu) + C_{00}^{\text{sat}}(\mu)$ to the unweighted points. This fit is performed iteratively with additional 3 sigma clipping until convergence. The “PTC turnoff” is defined as the signal level of the last unmasked point following the fit to only the exponential approximation, typically corresponding to the maximum of the variance curve and generally higher than the rolloff threshold. The “PTC rolloff” is the signal level where $C_{00}^{\text{prelim.}}(\mu)$ and $C_{00}^{\text{exp. approx}}$ deviate by more than 1%.

Final model fit—With the remaining points below the PTC rolloff, we fit the complete covariance expansion from P. Astier et al. (2019) (their Eq. 20):

$$C_{ij}^{\text{model}}(\mu) = \frac{\mu}{g} \left[\delta_{i0}\delta_{j0} + a_{ij}\mu g \right. \\ \left. + \frac{2}{3} ([\mathbf{a} \otimes \mathbf{a} + \mathbf{a}\mathbf{b}]_{ij}) (\mu g)^2 \right. \\ \left. + \frac{1}{6} (2\mathbf{a} \otimes \mathbf{a} \otimes \mathbf{a} + 5\mathbf{a} \otimes \mathbf{a}\mathbf{b})_{ij} (\mu g)^3 \right. \\ \left. + \mathcal{O}[(\mu g)^4] \right] + \frac{n_{ij}}{g^2}, \quad (11)$$

where a_{ij} (electrons)^{−1} represents the fractional change in pixel area per stored charge due to boundary motion from the brighter-fatter effect, b_{ij} (electrons)^{−1} captures weaker time-dependent corrections, g is the gain

(electrons/ADU), and n_{ij} is the constant noise contribution (electrons)² (e.g., $n_{00} \sim (5.7 \text{ electrons})^2$ on the diagonal, with off-diagonal terms determined by read electronics and overscan). The symbol $\mathbf{ab}$ denotes element-wise matrix multiplication, and $\otimes$ represents discrete convolution. In practice, the b_{ij} parameter is inferred to be just noise around a value of zero. The size of the covariance matrix ($n \times n$) is determined by the range of the brighter-fatter effect (see §4.8). Figure 12 shows an example fit for amplifier C03 on detector 94, illustrating the measured C_{00} points, the preliminary exponential approximation with saturation roll-off, and the final variance from Eq. 11.

Because the gains are fit independently per amplifier, small fitting residuals can leave flux discontinuities at amplifier boundaries within a detector. We add a subsequent step that measures the relative offsets between amplifiers in flat fields after gain correction and computes small adjustments to the gain to make the detector image continuous while preserving the average gain across the detector. We compute the average gain adjustment in flat fields across a range of signal levels, and use these to produce overall “adjusted gains”. The adjustments to the gain are always small (less than 0.03), and uncorrelated between detectors. These final gains are used for conversion of science images to electron units and construction of the variance plane for source detection.

The PTC calibration products are applied during instrument signature removal of science images in three ways. First, the gain converts pixel values from ADU to electrons. Second, the gain and read noise inform variance plane estimation. Third, the saturated pixel mask flag (denoted “SAT”) is set for all pixels with values greater than or equal to the PTC turnoff threshold.

4.8. Brighter-Fatter Effect

The brighter-fatter (BF) effect is the first influence the detector has on an incident signal (Figure 3). Accumulated electrons in a pixel repel incoming charges into neighboring pixels. Brighter point sources therefore appear wider in the recorded images. This electrostatic effect reduces to an effective change in pixel area and introduces correlations between nearby pixel values. The effect is intrinsic to thick CCDs and can significantly bias point-spread-function (PSF) photometry and shape measurement.

The BF effect has been extensively characterized for LSSTCam by laboratory and in-dome commissioning measurements (A. Broughton et al. 2024). The strength of the effect is anisotropic ($a_{01}/a_{10} \sim 3$), long-ranged (~ 10 px), and wavelength dependent. It is also about

Table 1. Parameter Priors for the Electrostatic Brighter-Fatter Effect (BFE) Model Fit

Parameter	Prior
Si thickness, t (μm)	100 (fixed)
Pixel size, p (μm)	10 (fixed)
Charge cloud height, z_q (μm)	$0 \leq z_q \leq t/2$
Charge cloud length, a (μm)	$10^{-5} \leq a \leq 0.35 p$
Charge cloud width, b (μm)	$10^{-5} \leq b \leq 0.35 p$
Potential well vertex height, horizontal z_{sh} (μm)	$0 \leq z_{sh} \leq t/2$
Potential well vertex height, vertical z_{sv} (μm)	$0 \leq z_{sv} \leq t/2$

NOTE—Priors on the fit parameters are informed by pixel-level simulations of the electric fields in these pixels (C. Lage et al. 2017). The potential well vertices refer to the heights above the front-side of the CCD to the saddlepoints in the electric field that defines the pixel boundaries in the vertical (0, 1) and horizontal (1, 0) directions. The electrostatic model assumes that the stored charge cloud in the sensor takes the form of a flat, rectangular sheet of charge with length (a) and width (b). The inequality constraints $z_{sh} > z_q$, $z_{sv} > z_q$ are enforced during optimization, which is true in the unsaturated signal level regime.

twice as strong in the E2V-family detectors as in ITL-family detectors.

Several BF correction methods have been proposed and implemented in the LSST Science Pipelines (D. Gruen et al. 2015; W. R. Coulton et al. 2018; P. Astier & N. Regnault 2023; A. Broughton et al. 2024), and they differ widely in their underlying assumptions. A. Broughton et al. (2024) found that modeling individual boundary shifts for every pixel (N, S, E, W) is required for adequate correction at the level needed for LSST Y10 science requirements. Flat-field pixel correlations can measure only the net per-pixel area change, so determining these boundary shifts requires additional electrostatic modeling. For processing raw images from DP2, we use the P. Astier & N. Regnault (2023) algorithm developed for Hyper Suprime-Cam, which derives from electrostatic physics and treats each pixel boundary individually.

Calibration of the BFE using this algorithm takes the form of modeling a charge cloud at the bottom of a pixel, computing how the effective areas of nearby pixels would be expected to change electrostatically in response. It adjusts some geometric parameters (summarized in Ta-




Figure 12. Example PTC fit for amplifier C03 on LSSTCam detector 94. Measured C_{00} covariance points are compared to the preliminary fit (exponential approximation plus saturation roll-off) and the final full-covariance model variance. Vertical markers indicate the PTC rolloff and turnoff thresholds.

ble 1), and solves Gauss’s law to compute the shift of pixel boundaries and the relative change in pixel area from each shift ($\mathbf{a}^N$, $\mathbf{a}^S$, $\mathbf{a}^E$, $\mathbf{a}^W$). We optimize these parameters to best fit the measured net pixel area changes ($\mathbf{a}$ from Eq. 11):

$$\mathbf{a} = \mathbf{a}^N + \mathbf{a}^S + \mathbf{a}^E + \mathbf{a}^W. \quad (12)$$

In principle, the BFE can vary amplifier-to-amplifier, but we have observed only noise-level variations in the

pixel distortion matrices across amplifiers and even detector families. For this reason, and to avoid discontinuous amplifier boundaries, the fit is performed at the detector level. All a_{ij} matrices from well-behaved amplifiers are 3σ -clipped elementwise and averaged into a detector-mean a_{ij} ; the scatter becomes the per-element uncertainty, and a single fit is run against this average pixel distortion matrix. The output calibration data

product is the set of model parameters and a model that can be evaluated per-detector.

Applying this calibration to remove the effect during ISR uses the boundary-shift matrices $\mathbf{a}^N$, $\mathbf{a}^S$, $\mathbf{a}^E$, and $\mathbf{a}^W$ to redistribute flux accordingly. We tile these into two, $(2n + 1) \times (2n + 1)$, convolution kernels K^H and K^V for the horizontal and vertical shifts, respectively; we apply sign flips and edge adjustments so each matrix sums to zero. These matrices need not be symmetric, which is more realistic per A. Broughton et al. (2024). Figure 13 shows an example pair of these kernels for amplifier C03 on detector 94, illustrating the dipole-like, charge-conserving structure that redistributes flux across the vertical and horizontal pixel boundaries.

For an image Q_{ij} in natural gain-corrected units of electrons, the amount of charge lost over the vertical and horizontal boundaries is:

$$\begin{aligned}\delta_{ij}^H &= \frac{1}{2} \sum_{kl} K_{kl}^H Q_{i-k, j-l}, \\ \delta_{ij}^V &= \frac{1}{2} \sum_{kl} K_{kl}^V Q_{i-k, j-l},\end{aligned}\quad (13)$$

The $1/2$ prefactor accounts for the average BFE assuming a constant rate of charge accumulation during an exposure. The corrected image is assembled by removing the charge each pixel lost across its N/E boundaries and adding what it gained from S/W neighbors, guaranteeing charge conservation by construction, $\sum_{ij} Q_{ij}^{\text{corr}} = \sum_{ij} Q_{ij}$. This correction does not require iterative convolutions. Unlike the W. R. Coulton et al. (2018) algorithm used for DP1 (Vera C. Rubin Observatory Team 2026), which requires multiple convolutions to reach convergence, we perform exactly two convolutions with the horizontal and vertical boundary-shift matrices.

4.9. Bias

Due to (CZW: reasons), there may be small static signals that appear in all exposures, regardless of exposure time. To remove this remnant signal that is not already removed by the previous calibrations, a combined bias frame is subtracted from the exposure. This combined bias is constructed from the per-pixel clipped mean of a large number of input zero-second bias frames. These input bias frames have all of the corrections discussed up until this point applied. For LSSTCam, these combined biases have 400 input exposures, which reduces the noise in the correction. CZW: quick study to see if we have any with fewer than 400 exposures.

4.10. Dark

Similar to the bias, we construct a combined dark frame from a large number of 30-second exposures, with




Figure 13. Example brighter-fatter correction kernels K^V (vertical) and K^H (horizontal) derived from the electrostatic BFE model for amplifier C03 on LSSTCam detector 94. The diverging color scale shows the fractional charge redistributed across pixel boundaries; each kernel sums to zero to conserve charge and exhibits the antisymmetric dipole structure that moves flux from the brighter pixel to its neighbors.

that time chosen to match the standard exposure time of a science exposure. As before, all of the previous calibrations are applied to these dark frames, including the bias correction. These exposures have an additional cosmic ray rejection step applied after the ISR processing, as these longer exposures may have more of these features than are found on the bias frames. The combined dark is again created from the per-pixel clipped mean, and for LSSTCam we have a target of 400 input frames. CZW: I know we don't always hit that, put words here.

4.11. Flats

Although flat field images are a common astronomical correction, the large field of view of LSSTCam make gen-

erating this calibration a particular challenge. In addition to ensuring consistent illumination across this field of view, there is an added complication that the ITL and E2V detectors have different responses as a function of wavelength. One final complication comes from the finite choices of LEDs in the flat field illumination system described above. We anticipate that we will be able to construct higher quality synthetic flats from sequences of laser-illuminated monochromatic flats in the future, for DP2 we have chosen to construct “anaglyph” flats for most bands.

As the name suggests, these anaglyph flats are constructed from two sets of flat field images that have each been observed with a single active LED illuminating the flat field screen. For the subsequent discussion, these two sets of input flat field images are referred to as the “red” and “blue” exposures, depending on the wavelength of the two LEDs within the filter bandpass that final flat field will be used for. These anaglyph flats have been constructed for each of the *grizy* filters. This process is unavailable for the *u* filter, as sufficiently bright far ultraviolet LEDs are not currently available, and so flats for this filter are constructed from only a single set flats (with an LED that is at the red end of the bandpass).

To begin the construction of new flats, the input exposures are processed through the ISR steps of that are applied prior to flat correction, including the bias, dark, crosstalk, gain application, linearization, CTI, and overscan (CZW: reorder to be in order). Statistics are measured on these processed exposures, determining the median value and dispersion on each detector. Normalization factors are then measured for the red and blue inputs separately based on these statistics, which should account for any differences in the absolute flux received in each exposure over all detectors. These normalization factors are used to properly scale the inputs before they are combined using a clipped mean per pixel into separate intermediate red and blue flats.

These intermediate flats include gradients related both to how well the flat field illumination and screen were oriented and due to the vignetting of the focal plane itself. Both types of gradients are measured next, so that contribution can be removed from the final flat. An initial six-node radial spline model is fit to the full binned focal plane image of the intermediate flats. This initial model is used to constrain both the centroid of the focal plane and the observed ratio of the ITL to E2V detectors, to account for their possibly different response. These parameters are then used for a 15-node radial spline model that is more careful about fitting the observed vignetting at high focal plane radii. This sec-

Band	Blue weight	Red weight
<i>g</i>	0.63	0.37
<i>r</i>	0.50	0.50
<i>i</i>	0.476	0.524
<i>z</i>	0.65	0.35
<i>y</i>	0.52	0.48

Table 2. Mixture weights for anaglyph flat construction.

ond fit uses a fixed ITL/E2V ratio measured during the first pass, to reduce the degrees of freedom in the final fit. An additional planar gradient is also included in this second fit, to account for simple illumination differences.

The measured gradients are then removed the intermediate flats. The planar gradients are simply divided from the per-detector intermediate flats. The radial gradient is scaled by a reference gradient to reduce the effects of the vignetting, and to match the vignetting observed in on-sky data (CZW: I need to explain reference gradient better). A final normalization factor is applied to scale the entire flat to match the intensity level observed within the central 25mm of the focal plane. With these gradient-corrected intermediate flats generated, a final flat is constructed by combining the red and blue intermediates with a filter-specific mixture weight, as listed in Table 4.11.

5. INSTRUMENT SIGNATURE REMOVAL FOR DP2 SCIENCE DATA

Figure 14 lists the ISR correction steps in the order they are applied to science images for DP2. This section integrates the calibration products described in §4 and summarizes how they compose during routine science processing. Corrections are applied in reverse chronological order relative to the photon transfer chain (§2.2 and Figure 3). Steps that undo readout-electronics effects therefore appear before steps that undo detector-level effects, and additive corrections precede multiplicative ones such as flat-fielding.

5.1. ISR Pipeline Overview

The LSSTCam has a particular photon transfer chain and therefore a particular ISR correction pipeline. The following outline follows Figure 14 step by step. Each item names the correction applied, and the reason it sits at this location in the processing chain.

1. **Convert integer to float.** Raw images arrive as integer arrays. This first step promotes pixel values to 64-bit floating point precision.
2. **Dither counts.** A uniform random offset in $[-0.5, 0.5)$ ADU is added to every pixel. This




Figure 14. Outline of the Instrument Signature Removal Pipeline (ISR) applied to LSSTCam science images for DP2. The steps in the pipeline follow a very specific sequence determined by the physics of the instrument signature and where it enters in the photon transfer chain. Blue boxes denote static ISR steps that apply pre-computed calibration inputs; orange boxes denote steps whose application depends on the science exposure itself (for example, overscan subtraction, dithering, or saturation-footprint growth); white boxes mark the raw input and post-ISR output. The dashed optional step applies pre-computed per-amplifier offsets when enabled.

breaks the correlation between true signal and integer quantization before continuous corrections are applied, reducing systematic bias when later steps interpolate, convolve, or divide the image. This step is essentially a marginalization step before a full low-signal, differential non-linearity correction can be measured and implemented.

- 3. Serial overscan correction.** An electronic, clock-injected offset plus smaller row-to-row variations are derived from the serial overscan and subtracted from the imaging pixels while the data are still in native analog-to-digital units (ADU). Overscan correction must precede gain conversion because the overscan bias is measured and removed in the same units as the raw readout.

- 4. Mask saturated and suspect pixels.** Pixels above the PTC turnoff saturation threshold (§4.7) get flagged as both saturated (“SAT”) and suspect (“SUSPECT”) in the pixel bit mask before further signal corrections. Early masking prevents corrupted or non-physical pixel values from biasing downstream estimates that assume a well-behaved signal.

- 5. Apply gain.** Per-amplifier gains derived from the PTC (§4.7) convert pixel values from ADU to electrons. Electron units are the native domain for deferred-charge modeling, brighter-fatter correction, and variance plane estimation, all of which were calibrated assuming a linear conversion from measured counts to stored charge.

6. **Crosstalk correction.** Electrical coupling between adjacent amplifier outputs is removed using the crosstalk matrix (§4.3). Crosstalk is applied after gain conversion so that coupled signals are corrected in electron units, consistent with how the matrix was measured.
7. **Parallel overscan correction.** Analogous to the serial overscan step, the parallel overscan region provides a column-direction estimate of readout offset that is subtracted from the imaging area. It follows crosstalk correction because amplifier coupling can inject signal into overscan regions, and removing crosstalk first yields a cleaner overscan bias estimate.
8. **Non-linearity correction.** The double-spline linearizer (§4.5) maps measured counts to a linearized electron count by subtracting absolute and amplifier-relative spline corrections. This step must precede deferred-charge correction because the deferred-charge model and its calibration were fit assuming that signal-dependent readout nonlinearities have already been removed (§4.6).
9. **Deferred charge correction.** Residual charge left in pixels during serial readout—from global charge transfer inefficiency, local electronic offset in the output amplifier, and serial charge traps on ITL sensors—is corrected using the parameters fit in §4.6. This step sits late in the readout-electronics portion of the chain because deferred charge is imprinted during serial transfer, after non-linearities in the output path have acted on the signal. We do not correct deferred charge on E2V sensors because the deferred charge level is already below the level defined by [Ž. Ivezić & The LSST Science Collaboration \(2018\)](#).
10. **E2V edge-bleed mask.** We flag the “SUSPECT” mask plane bit for pixels in E2V amplifiers that show uncorrectable charge spill from bright regions at the serial-register edge.
11. **ITL dip, SatSag, and edge-bleed masks.** We flag the “SUSPECT” mask plane bit for ITL sensors that exhibit a pixel column depression in response near certain amplifier boundaries (the “ITL dip”), a sag in apparent signal after saturation (“SatSag”), and edge bleeding analogous to the E2V case. These vendor-specific masks follow deferred-charge correction because the artifacts they mark are readout- and saturation-related structures that should be evaluated on linearized, CTI-corrected data.
12. **Bias correction.** A combined bias frame (§4.9) is subtracted to remove the fixed-pattern electronic offset that remains after overscan subtraction. Overscan correction handles the spatially static electronic structure imprinted in the images before charge readout.
13. **Dark correction.** The combined dark frame (§4.10), scaled by the science exposure time, is subtracted to remove thermal dark current accumulated during exposure integration. Dark correction follows bias subtraction because both are additive offsets, and the dark frame is constructed from bias-subtracted dark exposures.
14. **Defect masking.** Persistent bad pixels, columns, and other static defects identified during calibration (§4.4) are merged into the pixel mask. Applying the defect mask after deferred-charge, bias, and dark correction places it once readout-electronics effects and integration-period offsets have been removed, matching the processing state in which the defect catalog was built and ensuring that later defect footprint widening and interpolation operate on the full static bad-pixel map.
15. **Widen saturation footprints.** The saturation mask is expanded spatially to include pixels affected by charge spill and trails from saturated sources. This step follows defect masking so that all static bad regions are known before the saturation footprint is grown into neighboring pixels.
16. **Brighter-fatter correction.** The electrostatic brighter-fatter effect (§4.8) redistributes flux among neighboring pixels to undo charge-cloud repulsion in the silicon. In the photon transfer model (Figure 3), brighter-fatter is the earliest detector-level effect on incident photons, so its correction appears correspondingly late in ISR. It is applied after the image has been converted to linearized electron units and after additive calibration-frame corrections. The correction also needs to match the boundary-shifts, which were also calibrated from PTC covariances measured after these corrections.
17. **Variance plane estimation.** An expected variance is assigned to each pixel using the PTC-derived gain and read noise (§4.7). The variance plane is computed after gain application and brighter-fatter correction so that it reflects the electron units and pixel correlations relevant to source detection and weighting in downstream analyses.

At this stage the expected per-pixel variance follows the Poisson-plus-read-noise model

$$\sigma_{ij}^2 = \frac{F_{ij}}{g_a} + \left(\frac{\sigma_n(a)}{g_a} \right)^2, \quad (14)$$

where F_{ij} is the science image value in the current ISR units, g_a is the PTC gain for amplifier a (electrons per ADU), and $\sigma_n(a)$ is the read noise in electrons. For DP2 science processing the image is already in electron units, so $g_a = 1$ and Eq. 14 reduces to $\sigma_{ij}^2 = F_{ij} + \sigma_n(a)^2$. Pixels with non-positive variance are flagged in the mask.

Since the variance plane is used to estimate the uncertainty on just the charge accumulated in the image, it needs to be computed before flat-field correcting.

18. **Flat correction.** The normalized flat field (§4.11) is divided into the image to correct pixel-to-pixel variations in quantum efficiency and optical throughput. Flat-fielding is intentionally late in the pipeline: it is a multiplicative correction that must be applied only after all additive electronic and dark offsets have been removed, so that the flat field encodes relative sensitivity rather than residual bias structure.
19. **Interpolate masked pixels.** Pixels in the “SAT”, “BAD”, and “UNMASKEDNAN” mask planes are filled by interpolation from neighboring unmasked values; “SUSPECT” pixels are flagged but not interpolated by default. Interpolation is performed after flat-fielding so that both the science signal and the replacement values are in the same calibrated flux units expected by algorithms that assume a continuous image (for example, background estimation and certain astrometric measurements).
20. **Amplifier offsets (optional).** Small per-amplifier multiplicative adjustments derived from flat-field residuals may be applied as a final optional step to reduce residual flux discontinuities at amplifier boundaries. This step is optional because the PTC gain solution already includes average amplifier matching (§4.7); the offset correction addresses any remaining channel-to-channel offsets visible after the full ISR chain.

Together, these steps transform the raw integer read-out into a floating-point, full-calibrated image in electron units with an associated variance plane and pixel mask used in downstream data release processing tasks.

It is crucial for the ordering to match the physics of the detectors and the reverse chronological ordering of the photon transfer chain. Running the processing pipeline with misplaced steps would either apply a correction in the wrong physical units or leave residual structure that invalidates a downstream calibrations

5.2. Curated Calibrations and Header Provenance

Science exposures processed for DP2 are matched to a curated calibration collection (§4.1) appropriate to their calibration epoch (§3.3). The collection bundles the certified bias, dark, flat, defect mask, crosstalk matrix, linearizer, deferred-charge parameters, PTC table, and brighter-fatter distortion model, together with any auxiliary data products required for masking or optional corrections (§4.2). When the observing conditions, sequencer configuration, or calibration campaign change enough to warrant a new epoch, a new curated collection is produced and certified before it is attached to incoming science data.

Exposure header metadata records the calibration epoch, the identifier of the curated collection used, and a per-step record of which ISR corrections were applied or skipped. This provenance makes it possible to reproduce the processing of any released image, compare products generated across software or calibration versions, and audit whether optional steps such as amplifier-offset correction were enabled for a given dataset. The verification, acceptance, and distribution context for these metadata fields is described in §7.6.

6. DISCUSSION

Describe how calibrations/ISR have changed across different calibration epochs (sequencer changes).

List the set of criteria that prompt a new round of calibrations.

7. VERIFICATION, ACCEPTANCE, AND CERTIFICATION

Before calibrations are used for science processing, they first undergo an internal verification process to ensure they improve images that are processed with them. As these calibration products are used for all parts of the survey, approving updates is done via an internal review that is open to any member of the project, to ensure that all views are heard before processing uses them. This section details the steps used to verify, accept, and deploy calibration products for use as part of the survey. These steps do not preclude users from constructing and using their own set of calibrations, but they ensure that the default set of calibrations are suitable for use. Full details about the commands used for

Calibration Product	Quantity
Bias	Image mean Image stdev
Dark	Image stdev with CR rejection
CZW:	Add more examples

Table 3. Brief summary of verification tests.

these steps, and suggestions for a consistent collection naming system are presented in DMTN-222 (C. Z. Walters & E. S. Rykoff 2026).

7.1. Verification

After new calibrations have been constructed, they are used to process some set of exposures to show that they correct the various instrument features as expected. For simpler calibrations, such as bias and darks, it is usually sufficient to apply the combined calibration to their input bias and dark frames, to confirm that the residuals are consistent with only the read noise remaining. Other calibrations require more analysis, such as applying the brighter-fatter correction to a series of on-sky images and confirming that the distribution of source sizes is no longer a function of the source flux. Calibrations that do not pass some tests are not necessarily rejected outright, although those failures should be raised and considered during the acceptance process. This is particularly true of tests that examine a property at the amplifier segment level; rejecting a calibration due to a small number of problematic segments is impractical. Where calibration construction does not produce a result that passes all of the requirements, we can add to our defect list to exclude regions that have serious issues from further processing.

The current set of tests used, and the criteria used for acceptance are detailed in DMTN-101, with a brief summary listed in Table 7.1. It is likely that we will expand the list of verification tests calculated for calibrations as our familiarity with the camera and data improves.

7.2. Acceptance

Once a set of calibrations has been generated and verified, a meeting of the Telescope and Auxilliary Instrumentation Calibration Acceptance Board (TAXICAB) is called to decide if those calibrations should be deployed for use. The TAXICAB does not have a fixed membership, and welcomes any member of the project to join to offer opinions and comments on the calibrations presented. The goal of this board to ensure that the decision making for calibration approval is public and accessible, with the final decisions made by consensus agreement. Associated with each TAXICAB meet-

ing is a “hailing ticket,” which is used to provide a location for details and artifacts from the generation and verification process. These artifacts of the generation and verification process are presented to the TAXICAB meeting, where the attendees can decide whether or not the calibration is sufficient, and whether further tests are needed to determine if the calibration performs as expected. A calibration with failures of some of the verification tests is not necessarily excluded, and may be allowed for use based on the significance of the failures, with the TAXICAB making that final decision.

7.3. Certification

During the calibration construction process, the calibrations are generally certified for a particularly validity date range. This date range is allowed to be open ended in either direction, but in most cases we are constructing calibrations with a known start date (for instance, the start of observing with LSSTCam) and an open end date, allowing that calibration to be used until replaced by a future calibration. Certification is a butler database operation that adds the individual datasets that comprise the calibration to a second CALIBRATION collection that contains these validity date ranges. This process allows the butler system to identify which calibration datasets should be supplied for processing raw images, with the first entry in a list of CALIBRATION collections that has a validity range that includes the observation date and time of those individual raw images. By constructing CHAINED collections that point to multiple CALIBRATION collections, each containing some set of calibration products and validity ranges, users can ensure that the butler will supply the correct calibrations for arbitrary input exposures without manual selection. As the butler searches the collection chains for the first appropriate calibration, by prepending calibration updates to the start of the chain, those updates will be found and returned before calibrations found in later collections, allowing calibrations with infinite future validity to be replaced without changing their validity date ranges.

7.4. Distribution

Adding a new CALIBRATION collection to the main top-level CHAINED collection holding all of the calibration data for LSSTCam is done only after acceptance of the calibration products it contains by the TAXICAB. This is the first step of the distribution process, and allows the butler repository where the calibration was constructed to be available for use. However, there are multiple butler repositories that are used for production, and so the calibrations and associated metadata

must be transferred to all of these locations as well. The CALIBRATION collections holding the data for a given TAXICAB ticket are next exported to create a calibration archive containing a copy of each file that make up the associated calibration products, with a YAML dump of the necessary butler metadata to allow this exported archive to be imported to a different repository. Consistent with the other data curation processes used to synchronize data between the data facilities that will process and serve Rubin data (CZW: asdfasdf, this is a bad sentence), we use the Rucio ([Rucio](#) [????](#)) system to register and transfer the calibration files. Rucio is not instantaneous, and must wait until the files have been copied to the storage local to each data facility. Once that copy is complete, the YAML dump of the butler metadata is used to import these calibration products into the local butler repository for use.

7.5. Data Monitoring

One additional advantage of the verification system is that it provides a way to monitor the state of the camera, and the continuing appropriateness of the calibration products. By running the verification tasks on new calibration exposures as they are taken, we can check how well the existing calibrations correct these new exposures. Although not running during the DP2 observations and analysis, this verification will be added to the processing performed at the summit, with the metrics calculated from this processing passed to the Sasquatch and Chronograf monitoring system that is used to track other performance quantities. This monitoring will allow the health of the camera to be tracked, and for the generation of new calibrations to be prompted by any observed changes in the metrics over time.

7.6. Data products, Distribution, and User Access

Describe the output schema, distribution process (exporting/importing, rucio, TAXICAB tracking), and user access.

We enumerate the primary data products (calibration frames, ISR outputs, masks, and metadata tables), describe naming/versioning conventions, and summarize the provenance tracked for each product (inputs, software versions, and configuration).

Briefly describe how users discover and retrieve products (e.g., via a data butler/repository interface) and recommended query patterns.

Be sure to mention that header metadata contains information about whether a particular ISR step was applied/not applied and which calibration dataset was used.

8. LIMITATIONS & KNOWN ISSUES

The ISR pipeline developed for LSSTCam removes the dominant instrumental signatures, but several effects remain partially corrected, unmodeled, or unmasked. The paragraphs below summarize known limitations (instrumental systematics not fully corrected, edge cases, validity constraints, etc.) for interpreting DP2 data, our recommendations for handling them in the context of DP2 science, and plans for near-term improvements expected before Data Release 1 (DR1).

Flares—Some science images contain localized, bright flare-like structures. They appear as elongated trails that begin at a detector edge and fade with a wavelength-dependent power-law profile consistent with attenuation of light in silicon at that wavelength. The light is likely scattered from bright stars: focused starlight may reflect from the beveled edge of the silicon mount and enter a neighboring detector through its side, although this mechanism is not yet confirmed. Similar structures appear in some flat-field images as many small edge flares (referred to as a “comb” pattern), possibly linked to the rough-cut silicon edge. So far, these features have been seen only in y-band, with both focused and unfocused illumination, and they are transient, consistent with a dependence on the chance alignment of bright stars with the focal-plane array. They are not yet fully characterized or masked in DP2 processing. Science users should treat anomalously bright detections/extended objects near detector edges with caution in y-band single-visit and difference images, and verify suspect sources by visual inspection.

Residual non-linearity wiggles—The double-spline non-linearity model (§4.5) removes the dominant signal-dependent departure from linear response, but small quasi-oscillatory residuals remain as a function of signal level. They appear as low-amplitude, signal-dependent structure in linearized images. These “wiggles” are irregular, clearly observed beyond noise fluctuations, and they do not correlate with measured properties of the input flat-field images (e.g. time, temperature) or with channel on a detector or across the focal plane. These residuals are well below current science requirements and are left uncorrected in DP2, but further characterization is in progress.

Differential non-linearity (DNL)—The DP2 linearizer corrects average amplifier non-linearity but does not yet apply a full low-signal differential non-linearity correction. This limitation reflects the difficulty of delivering flat-field illumination below 100 ADU/pixel with the indome calibration system. Users should exercise caution

when relying on photometric or astrometric measurements at low signal levels, in particular when relying on u-band sky-background estimates from later processing steps, where backgrounds are typically at or below 100 ADU/pixel. Count dithering in ISR (§5.1) provides interim mitigation until a dedicated DNL correction is measured and deployed.

Detector vignetting models—Flat-field calibrations measure relative pixel and filter quantum efficiencies and throughputs under uniform illumination, but they require an accurate model of optical vignetting across the focal plane. Improved vignetting models, informed by larger calibration datasets and cross-checks against on-sky stellar photometry, are planned for future calibration epochs to reduce focal-plane non-uniformity.

Color-dependent brighter-fatter corrections—The brighter-fatter effect is color dependent because its amplitude depends on photon conversion depth in silicon and on the path length over which the transverse electric field acts on drifting electrons. Laboratory and commissioning measurements show that the effect is chromatic ($a_{00}^{y\text{-band}} = 0.94 \times a_{00}^{r\text{-band}}$, see A. Broughton et al. 2024). The electrostatic brighter-fatter model used in DP2 (§4.8) is fit from PTC covariances and applies a single wavelength-averaged correction per detector, derived in the band of the PTC input flats. Because the effect is strongest near the bottom of the potential well and the front of the chip, calibrations for u through i bands are approximately consistent as most light in those bands converts to photoelectrons within the first few microns. Separate calibrations are appropriate for z- and y-band light, which typically convert far deeper into the silicon. Boundary-shift calibrations can be color-corrected by assuming a flat incident spectral energy distribution in each filter band and a depth-conversion model in silicon. These corrections were not enabled in DP2 because validation was incomplete, but they are expected for DR1. Until then, PSF photometry and shape measurements of bright sources, especially in redder bands, may retain small flux- and filter-dependent biases.

Fringing—Fringing is an interference pattern produced when narrow-band emission reflects within the thin, back-illuminated silicon of LSSTCam CCDs. It is driven primarily by night-sky emission lines and becomes appreciable in the z and y bands, where the silicon is increasingly transmissive and OH and other atmospheric lines contribute strongly to the background. In-dome flat-field exposures (§4.11) imprint and partially divide out a fringe pattern, but the fringing amplitude and phase in science images depend on the time-varying on-sky spectrum and therefore are not fully canceled by

flat-fielding alone. The DP2 ISR pipeline (§5) does not include a dedicated fringe-subtraction step. Residual fringe structure is visible at low surface brightness in z- and y-band visit images and can bias wide-aperture photometry. Fringing models developed by Z. Guo et al. (2023) and corresponding corrections are under development and expected for DR1; until then, users analyzing faint extended emission or precision surface-brightness measurements in redder bands should inspect visit images for quasi-periodic fringe patterns and treat affected regions with appropriate caution.

Masking—Several limitations for the defects exist for DP2:

- vampire pixels
- edge bleeds

Further developments will provide better support and masking for LSST data releases.

9. CONCLUSION

This paper documents the calibration-production algorithms, instrument signature removal (ISR) workflow, and quality-assurance framework used to generate DP2 calibration products for LSSTCam. DP2 is the first end-to-end demonstration of data processing with on-sky LSSTCam dataset, and it establishes a reproducible reference for the first stage of processing in Data Release Processing (DRP) pipeline.

We organize calibration production and ISR around the electronic detector model introduced in §2.2, where each instrument artifact get its own associated calibration dataset and correction algorithm. Calibrations are built on top of other calibrations (summarized in Figure 5), requiring intermediate applications of ISR (summarized in Figure 14). In general, the pipelines are derived and applied in the reverse chronological order of the photon transfer chain. For DP2, we describe production algorithms and ISR for crosstalk, defect masking, non-linearity, deferred charge (CTI), the photon transfer curve (PTC), the brighter-fatter effect, and the classical bias, dark, and flat-field calibrations (§4), and we integrate their application in the science ISR pipeline (Figure 14, §5).

Where detailed validation is already in place, DP2 calibrations meet or approach the requirements needed for commissioning-era science. These products are accompanied by a verification, acceptance, and certification workflow (§7) intended to make calibration quality auditable, and we record the relevant provenance in exposure dataset metadata for the science community (§5.2).

DP2 nevertheless reflects commissioning-era conditions: calibration data were collected before dome completion and full thermal stabilization (§1), detector operating modes and sequencer configurations evolved across epochs (§6), and calibrations and correction tools for some artifacts remain under active development (§8). Users should treat residual focal-plane non-uniformities, epoch-dependent systematics, and incomplete coverage effects as known limitations when interpreting DP2 data products and deciding paths/configurations for science analyses.

Looking toward Data Release 1 (DR1), we expect iterative improvements in three areas: *algorithms* (refined CTI and trap modeling, expanded color-dependent brighter-fatter treatment), *operations* (more stable dome and calibration-system performance, clearer calibration-epoch boundaries, and expanded automation through verification and acceptance pipelines), and *diagnostics* (trending analyses, regression and alerts for calibration changes against baselines, and quality metrics on science-images for downstream pipeline acceptance criteria). The DP2 workflow documented here is the baseline from which those improvements will be measured; maintaining versioned, reproducible calibration production will remain essential as LSSTCam tran-

sitions from data previews during commissioning to data releases over the full 10-year survey.

ACKNOWLEDGMENTS

This material is based upon work supported in part by the National Science Foundation through Cooperative Agreements AST-1258333 and AST-2241526 and Cooperative Support Agreements AST-1202910 and AST-2211468 managed by the Association of Universities for Research in Astronomy (AURA), and the Department of Energy under Contract No. DE-AC02-76SF00515 with the SLAC National Accelerator Laboratory managed by Stanford University. Additional Rubin Observatory funding comes from private donations, grants to universities, and in-kind support from LSST-DA Institutional Members.

Facility: Rubin

Software: LSST Science Pipelines

REFERENCES

- Aihara, H., AlSayyad, Y., Ando, M., et al. 2022, PASJ, 74, 247, doi: [10.1093/pasj/psab122](https://doi.org/10.1093/pasj/psab122)
- Astier, P., Antilogus, P., Juramy, C., et al. 2019, A&A, 629, A36, doi: [10.1051/0004-6361/201935508](https://doi.org/10.1051/0004-6361/201935508)
- Astier, P., & Regnault, N. 2023, A&A, 670, A118, doi: [10.1051/0004-6361/202245748](https://doi.org/10.1051/0004-6361/202245748)
- Bosch, J., Armstrong, R., Bickerton, S., et al. 2018, PASJ, 70, S5, doi: [10.1093/pasj/psx080](https://doi.org/10.1093/pasj/psx080)
- Broughton, A., Utsumi, Y., Plazas Malagón, A. A., et al. 2024, PASP, 136, 045003, doi: [10.1088/1538-3873/ad3aa2](https://doi.org/10.1088/1538-3873/ad3aa2)
- Collaboration, T. D. E. S. 2005, arXiv, doi: [10.48550/arxiv.astro-ph/0510346](https://doi.org/10.48550/arxiv.astro-ph/0510346)
- Coulton, W. R., Armstrong, R., Smith, K. M., Lupton, R. H., & Spergel, D. N. 2018, AJ, 155, 258, doi: [10.3847/1538-3881/aac08d](https://doi.org/10.3847/1538-3881/aac08d)
- Diehl, T., & Collaboration, F. T. D. E. S. 2012, Physics Procedia, 37, 1332, doi: [10.1016/j.phpro.2012.02.472](https://doi.org/10.1016/j.phpro.2012.02.472)
- Esteves, J. H., Utsumi, Y., Snyder, A., et al. 2023, PASP, 135, 115003, doi: [10.1088/1538-3873/ad0a73](https://doi.org/10.1088/1538-3873/ad0a73)
- Gruen, D., Bernstein, G. M., Jarvis, M., et al. 2015, Journal of Instrumentation, 10, C05032, doi: [10.1088/1748-0221/10/05/C05032](https://doi.org/10.1088/1748-0221/10/05/C05032)
- Guo, Z., Walter, C. W., Lage, C., Lupton, R. H., & LSST Dark Energy Science Collaboration. 2023, PASP, 135, 034503, doi: [10.1088/1538-3873/acbe67](https://doi.org/10.1088/1538-3873/acbe67)
- Ivezić, Ž., & The LSST Science Collaboration. 2018, LSST Science Requirements Document, Project Controlled Document LPM-17, NSF-DOE Vera C. Rubin Observatory. <https://ls.st/LPM-17>
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, ApJ, 873, 111, doi: [10.3847/1538-4357/ab042c](https://doi.org/10.3847/1538-4357/ab042c)
- Lage, C., Bradshaw, A., & Tyson, J. 2017, Journal of Instrumentation, 12, C03091, doi: [10.1088/1748-0221/12/03/C03091](https://doi.org/10.1088/1748-0221/12/03/C03091)
- Miyazaki, S., Komiyama, Y., Nakaya, H., et al. 2012, Ground-based and Airborne Instrumentation for Astronomy IV, 84460Z, doi: [10.1117/12.926844](https://doi.org/10.1117/12.926844)
- Morganson, E., Gruendl, R. A., Menanteau, F., et al. 2018, arXiv, 130, 074501, doi: [10.1088/1538-3873/aab4ef](https://doi.org/10.1088/1538-3873/aab4ef)
- Roodman, A., Rasmussen, A., Bradshaw, A., et al. 2024, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 13096, Ground-based and Airborne Instrumentation for Astronomy X, ed. J. J. Bryant, K. Motohara, & J. R. D. Vernet, 130961S, doi: [10.1117/12.3019698](https://doi.org/10.1117/12.3019698)

- Rubin Observatory Science Pipelines Developers, AlSayyad, Y., Axelrod, T. S., et al. 2026, The LSST Science Pipelines Software: Optical Survey Pipeline Reduction and Analysis Environment, Project Science Technical Note PSTN-019, NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2570545](https://doi.org/10.71929/rubin/2570545)
- Rucio. ????, Rucio Distributed Data Management Documentation, <http://rucio.cern.ch/>
- SLAC National Accelerator Laboratory, & NSF-DOE Vera C. Rubin Observatory. 2025, The LSST Camera (LSSTCam), SLAC National Accelerator Laboratory (SLAC), Menlo Park, CA (United States), doi: [10.71929/rubin/2571927](https://doi.org/10.71929/rubin/2571927)
- Snyder, A., Roodman, A., Utsumi, Y., et al. 2021, Journal of Astronomical Telescopes, Instruments, and Systems, 7, 048002, doi: [10.1117/1.JATIS.7.4.048002](https://doi.org/10.1117/1.JATIS.7.4.048002)
- Utsumi, Y., Antilogus, P., Astier, P., et al. 2024, in Ground-based and Airborne Instrumentation for Astronomy X, Vol. 13103, 131030W, doi: [10.1117/12.3019117](https://doi.org/10.1117/12.3019117)
- Vera C. Rubin Observatory Team. 2026, The Vera C. Rubin Observatory Data Preview 1, Technical Note RTN-095, NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2570536](https://doi.org/10.71929/rubin/2570536)
- Waters, C. Z., & Rykoff, E. S. 2026, Calibration Generation, Verification, Acceptance, and Certification, Data Management Technical Note DMTN-222, NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2586498](https://doi.org/10.71929/rubin/2586498)