

Image Calibration and Instrument Signal Removal for the First Year of Rubin LSSTCam

ALEXANDER BROUGHTON ¹, AGNÈS FERTÉ ², AND CHRISTOPHER Z. WATERS ³

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*

(Dated: June 8, 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.

1. INTRODUCTION

The Vera C. Rubin Observatory will conduct the 10-year Legacy Survey of Space and Time (LSST, [Ivezić et al. 2019](#)), imaging roughly 18,000 deg² of the southern sky in six broadband, optical–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 CCDs from two vendors (Teledyne e2v and International Technologies 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 (dome flats, LEDs, monochromator, 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 imprinted 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 (NSF-DOE Vera C. Rubin Observatory 2025) 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) (NSF-DOE Vera C. Rubin Observatory Team et al. 2025) 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. 2017) 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. §3 introduces the observatory hardware and instrumentation that supplied the inputs for DP2 calibrations. §4 presents the physical detector model and the calibration and ISR pipeline, including algorithms and data products used in production. §5 discusses calibration-epoch differences and residual systematics relevant to DP2. §6 describes

how calibration products are verified, accepted, certified, and deployed. §8 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 (<https://lsstcam.lsst.io/>) is the largest digital camera ever built to date. Its design and the technological progress it required, were developed in order to meet the requirements set by the ambitious goals of LSST, as described in [add citation to LSST Camera Verification Testing and Characterization and other SPIE papers](#). 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). We summarize here key points about LSSTCam and properties of its CCDs relevant for LSSTCam calibration in Rubin Data Management (DM).

2.1. LSSTCam CCDs

LSSTCam, is composed of 189 CCDs organized in rafts of 9 CCDs (3x3), in a focal plane mosaic presented in CTN-001 [add citation](#). 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 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 [add ITL case](#). 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 4.1.6 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

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. At Calibpalooza 2022 ([is their proceeding or document we can refer to?](#))




Figure 1. Here is an exposure before and after ISR.



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.

and during following years, we have indeed 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.](#)

The ISR pipeline implemented in the Rubin data processing aims at correcting each of these effects in the order opposite to which they happen.

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

Analyzed with updated sequencer (s3v2).




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.

to do: insert map as a percentage of used pixels: dead, bad and the dead ccd are in one color, color for CCDs worst than the median and green for best CCDs. Effective area: vignetted masked + defects.

3.2. Calibration System (CalSys)

Provide descriptions of auxiliary telescope/CBP/flat field illuminator/LEDs/monochromator and any instrumentation used for input datasets. Describe any system-level constraints that influenced calibration strategy.

Describe the LSST Camera and any technical limitations (bad detectors/amps). Introduce and define the sequencer.

3.3. Data Selection

Describe the data selection process for the DP2 dataset.

4. DATA PROCESSING OVERVIEW

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.

4.1. Calibration & ISR Pipeline

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. J. Roodman et al. \(2024\)](#) gives a detailed account of system characterization work and results from these campaigns.

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 4 summarizes calibration-product dependencies, and Figure 5 lists the ISR correction tasks in the order they are applied to science images.




Figure 4. 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.

4.1.1. Curated Calibrations

4.1.2. Auxiliary Data Products

4.1.3. Crosstalk

Describe how the crosstalk matrices are constructed.

4.1.4. Defects

We describe how defect masks are generated (bad pixels/columns, saturation trails, and other artifacts), the thresholds used, and how masks are combined across exposures.

4.1.5. Non-Linearity

We describe the linearity characterization procedure and the model used to correct detector response as a function of signal level. We specify the valid range and how extrapolation (if any) is treated.

4.1.6. 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 trap 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 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 4. 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 flat-field illumination sequence as the linearity dataset (§4.1.5), estimating the residual charge in electrons left in the overscan region 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 values 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 (§4.1.7), 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 §4.1.7) 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




Figure 5. 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.

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 6. 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 in 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 surviving after sigma-clipping a 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 only the local offset parameters, then the global CTI parameter, and then the trap capture curve from residuals between




Figure 6. 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 §4.1.7) 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.

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).

LSST Science Requirement 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 7 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 like an overall offset and 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 had small trap size that was difficult to measure beyond noise fluctuations and already negligible enough




Figure 7. Distribution of inferred global serial CTI across all science sensor amplifiers on the LSSTCam focal plane. *Blue*—Teledyne e2v (E2V) detectors; *Orange*—International Technologies Laboratory (ITL) Detectors. The fit to the deferred charge model allows global CTI to vary up to 1.0×10^{-5} . TODO: Add CTI-corrected versions.

to not need 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 produce are noisy and reveal no statistically significant charge residuals, so we do not apply any parallel deferred charge correction to science images.

4.1.7. 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.

The photon transfer curve used to calibrate LSSTCam for DP2 is measured from 543×2 flat field exposures at evenly spaced illumination levels as low 100 ADU and reaching at least 10% beyond saturation. At each illumination level, a pair of flats was acquired by varying both LED output and exposure time to achieve the tar-

get signal. We measure the covariance between pixels on the difference image of the two exposures to remove variations sourced by static, fixed-pattern artifacts.

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:

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

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

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

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:

$$C_{00}^{\text{sat}}(\mu) = \begin{cases} -e^{-(\mu-\mu_0)/\tau} & \mu \geq \mu_0 \\ 0 & \mu < \mu_0 \end{cases}, \quad (3)$$

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 fit, 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):

$$\begin{aligned} C_{ij}^{\text{model}}(\mu) = & \frac{\mu}{g} [\delta_{i0}\delta_{j0} + a_{ij}\mu g \\ & + \frac{2}{3}([\mathbf{a} \otimes \mathbf{a} + \mathbf{ab}]_{ij})(\mu g)^2 \\ & + \frac{1}{6}(2\mathbf{a} \otimes \mathbf{a} \otimes \mathbf{a} + 5\mathbf{a} \otimes \mathbf{ab})_{ij}(\mu g)^3 \\ & + \mathcal{O}[(\mu g)^4]] + \frac{n_{ij}}{g^2}, \end{aligned} \quad (4)$$

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, 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 §4.1.8). Figure 8 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. 4.

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 at 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.

4.1.8. 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 effective change in pixel area and introduces correlations between nearby pixel values. The effect is intrinsic to thick CCDs and can significantly bias PSF photometry and shape measurement.

The BF effect has been extensively characterized for LSSTCam by laboratory and in-dome commissioning




Figure 8. 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.

measurements (A. Broughton et al. 2024). The strength of the effect is anisotropic ($a_{01}/a_{10} \sim 3$), long-ranged ($\sim 10\text{px}$), and wavelength dependent. It is also about 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) algo-

Table 1. Parameter Priors for the Electrostatic 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.

rithm 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 Table 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 best-fit the measured net pixel area changes ($\mathbf{a}$ from Eq. 4):

$$\mathbf{a} = \mathbf{a}^N + \mathbf{a}^S + \mathbf{a}^E + \mathbf{a}^W. \quad (5)$$

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 aver-

age 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 9 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:

$$\begin{aligned} \delta_{ij}^H &= \frac{1}{2} \sum_{kl} K_{kl}^H Q_{i-k, j-l}, \\ \delta_{ij}^V &= \frac{1}{2} \sum_{kl} K_{kl}^V Q_{i-k, j-l}, \end{aligned} \quad (6)$$

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 (NSF-DOE Vera C. Rubin Observatory Team et al. 2025), which requires multiple convolutions to reach convergence, we perform exactly two convolutions with the horizontal and vertical boundary-shift matrices.

4.1.9. Bias

We describe how master biases are constructed (over-scan handling, trimming, combination statistic, masking) and the metrics used to monitor time stability.

4.1.10. Dark

We describe how master darks are produced (scaling by exposure time, cosmic-ray rejection, hot pixel treatment) and how residual structures are assessed.

4.1.11. Flats

We describe how flat-field calibrations are created (illumination correction strategy, normalization conventions, handling fringing/stray light) and how uniformity is quantified.




Figure 9. 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.

5. DISCUSSION

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

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

6. VERIFICATION, ACCEPTANCE, AND CERTIFICATION

This section defines the quality metrics for calibration products and ISR outputs. Here we define the criteria used to promote products to “accepted” and “certified”.

6.1. 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 what calibration dataset was used.

6.2. Verification

We list (possibly in a table) the automated verification tests for each data product (sanity checks, analysis plots, detailed plots, distributional checks, and science test images) and the thresholds that determine acceptance/rejection of different calibration data products.

6.3. Acceptance

This section will describe the calibration acceptance process (TAXICAB).

Describe how verified products are compared against earlier baselines to reveal regressions or deviations, and how acceptance decisions are documented.

6.4. Certification

Describe the final certification step in the context of the data “butler” and file system storage.

6.5. Data Monitoring

Discuss how we monitor incoming streams of new data for quality (e.g. rubinTV, plot navigator, trend monitoring with chronograph).

7. LIMITATIONS & KNOWN ISSUES

Summarize known limitations for DP2 (instrumental systematics not fully corrected, edge cases, missing calibrations, validity constraints) and near-term roadmap items.

8. CONCLUSION

We summarize the workflow and the key validation steps that ensure stable, reproducible calibration production and ISR. We also outline planned improvements (algorithmic updates, expanded diagnostics, and operational automation) and the expected impact on data quality.

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, Publications of the Astronomical Society of Japan, 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. 2017, Publications of the Astronomical Society of Japan, 70, S5, doi: [10.1093/pasj/psx080](https://doi.org/10.1093/pasj/psx080)
- Broughton, A., Utsumi, Y., Malagón, A. A. P., et al. 2024, Publications of the Astronomical Society of the Pacific, 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, Publications of the Astronomical Society of the Pacific, 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)
- Ivezić, ., Kahn, S. M., Tyson, J. A., et al. 2019, The Astrophysical Journal, 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)
- NSF-DOE Vera C. Rubin Observatory. 2025, PSTN-019: The LSST Science Pipelines Software: Optical Survey Pipeline Reduction and Analysis Environment,, NSF-DOE Vera C. Rubin Observatory Technical Report doi: [10.71929/RUBIN/2570545](https://doi.org/10.71929/RUBIN/2570545)
- NSF-DOE Vera C. Rubin Observatory Team, Acero-Cuellar, T., Acosta, E., et al. 2025, RTN-095: The Vera C. Rubin Observatory Data Preview 1,, NSF-DOE Vera C. Rubin Observatory Technical Report doi: [10.71929/RUBIN/2570536](https://doi.org/10.71929/RUBIN/2570536)
- Roodman, A. J., Rasmussen, A., Bradshaw, A. K., et al. 2024, Ground-based and Airborne Instrumentation for Astronomy X, 90, doi: [10.1117/12.3019698](https://doi.org/10.1117/12.3019698)
- 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)