

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

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 26, 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, [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 (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. §2 introduces the LSST Camera and the CCD readout architecture rel-

evant to calibration. §3 describes the DP2 calibration datasets, the calibration system, and the criteria used to select input exposures. §?? presents the data processing overview and the physical detector model that organizes calibration production and ISR. §4 describes calibration product production and application, including algorithms and data products used in production. §5 integrates how those products compose during instrument signature removal for DP2 science data. §6 discusses calibration-epoch differences and residual systematics relevant to DP2. §7 describes how calibration products are verified, accepted, certified, and deployed, and §8 summarizes known limitations for DP2. §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 (<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




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

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.

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. Figure 3 summarizes the instrumental effects identified in the photon transfer chain and their order.

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?](#)) 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

Describe the filters.

Describe the sequencer

Analyzed with updated sequencer (s3v2).

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)

LSST uses four calibration systems to construct calibrations. Three of which (the flat-field illuminator, the collimated beam projector, and the Auxiliary telescope)




Figure pending

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.

will be used continuously throughout the LSST, while one calibration system (the Camera Calibration Optical Bench) was used to characterize and qualify the LSST-Cam CCDs before integration to the Simonyi Survey Telescope at Rubin Observatory.

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) (?). The CCOB-wide beam is designed to project a uniform monochromatic light source onto the focal plane. Light is produced by one of six LEDs, each corresponding to a central wavelength in the LSST filter bandpass, is controlled by a custom LED driver and fed into an integrating sphere. The light from the integrating sphere is projected onto the focal plane with a minimal ND filter (10%) attached to a C-mount lens with an adjustable f/stop. The CCOB-wide beam was used in run 6 and run 7.

I mention run 6 and run 7 here - it is not properly introduced. Is it within scope to mention the EO-testing run at the Rubin Observatory L3 CR?

One difference between run 6 and run 7 was that the f/stop of the lens was changed from 2.6 to 1.6 (fully open). This was done to try to reduce the effect of the ‘weather’ and the ‘CMB pattern’ two effects that we found in Run 6 at SLAC, and were found to be due to our projection setup (?). While changing the f/stop did reduce the weather pattern, it also caused a much steeper illumination roll-off across the focal plane. We

evaluated the weather pattern and illumination roll-off relative to Run 6.

To reduce the effects of the ‘weather’ and ‘CMB’ but retain uniform illumination across the focal plane, we installed a Lambertian diffuser to the CCOB-wide beam for Run 7. The diffuser used is a 60° diffusing angle unmounted sheet from Edmunds Optics. We found that the diffuser greatly reduced the ‘weather’, eliminated the CMB pattern, more uniformly illuminated the focal plane, with a penalty of decreasing the overall illumination by roughly 35%. The diffuser was installed for all run 7 record runs using the CCOB Wide Beam.

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. The Auxiliary Telescope

The Vera C. Rubin Observatory’s Auxiliary Telescope (AuxTel) is an independent 1.2-meter telescope that incorporates both spectroscopic and imaging capabilities to support the LSST **Cite an AuxTel overview paper**. Situated on Calibration Hill, adjacent to the primary observing facility at Cerro Pachón, Chile, AuxTel was originally constructed and operated as the Calypso Telescope at the Kitt Peak National Observatory before being acquired by the project in 2008 and subsequently refurbished. As a core component of the broader Rubin calibration system infrastructure, AuxTel acts as a fully automated, independent facility developed to measure and characterize the wavelength-dependent atmospheric transmission profile above the site in parallel with pri-




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.

mary survey activities. Because localized atmospheric absorption represents a primary driver of overall photometric error, high-cadence temporal and spatial characterization of the atmosphere is essential to satisfy the strict photometric precision and accuracy constraints mandated by the LSST. Target selection and operational scheduling for this facility are handled completely autonomously by the AuxTel Scheduler (ATScheduler), a specialized software element designed to leverage both historical data and predictive footprints of the Simonyi Survey Telescope to optimize observation coordination across the observatory **Cite a Scheduler technical paper**.

To accomplish these calibration requirements, the facility relies on a dedicated instrumental payload installed at one of the telescope’s Nasmyth ports. This configuration is known as the LSST Atmospheric Transmission and Slitless Spectrograph (LATISS) **Cite a LATISS paper**. In its primary mode, LATISS utilizes a transmission hologram dispersing element optimized for slitless spectrophotometry within a convergent telescope beam. The orientation of the holographic dispersing element is such that the photometry of the stars that produce spectra can be simultaneously be captured

using the single $4K \times 4K$ CCD detector. This detector is identical in fabrication and performance characteristics to those integrated within the primary LSST Camera (LSSTCam). Under standard operations, projected exposure times for LATISS range from 30 to 60 seconds, depending on the target’s baseline brightness, local cloud attenuation, and the specific instrument configuration.

The data products generated during operation of AuxTel and LATISS fulfill a dual-purpose mandate spanning direct science calibration and system-wide operational readiness. By strategically targeting standard stars—with explicit priority given to sources cataloged within the Hubble CALSPEC and Gaia satellite data releases, the system records spectral shifts as a function of changing airmass and time to extract the true localized atmospheric transmission function. The resulting spectroscopic data products allow for the empirical decomposition and monitoring of the five primary physical constituents governing atmospheric throughput: water vapor (H_2O), molecular oxygen (O_2), ozone (O_3), Rayleigh scattering, and airborne aerosols **This would be a great place to add a citation, if we want to keep this brief discussion about the aerosols.**

Ultimately, these measurements deliver a tailored atmospheric transmission estimation mapped directly to each distinct observation executed by the primary survey telescope.

Commissioning campaigns tracking over 130,000 spectra up to mid-2026 demonstrate that fully capturing local atmospheric variability does not require strict, real-time pointing coordination with the main telescope; instead, it is sufficient for AuxTel to sample standard stars within a relaxed 5-minute temporal window and a 60° spatial radius of the Simonyi Survey Telescope footprint. **Cite the internal AuxTel report, if necessary. Also could cut this for brevity if needed..**

Beyond direct calibration products, the AuxTel facility serves as a vital real-world pathfinder test stand for the entire observatory’s software environment, training personnel and exercising 24-hour operational workflows. It has been instrumental in validating the Feature-Based Scheduler for autonomous target selection, deploying Service Abstraction Layer (SAL) data acquisition scripts, and refining the LSST Operations and Visualization Environment (LOVE) interface under a continuous observing cadence. Furthermore, dedicated imaging campaigns run on AuxTel generate live data streams used to test data release production, exercise pipeline workflows, and validate the prompt processing frameworks necessary for real-time alert generation and distribution well before the main telescope goes on-sky. **Not sure if this is in scope but I will leave it here, at your discretion whether or not to prune**

3.2.3. Collimated Beam Projector

To achieve the scientific objectives of the LSST, milli-magnitude level precision of photometric calibration is necessary. **Cite maybe?.** To meet these stringent requirements, the Collimated Beam Projector (CBP) was developed as a dedicated, in-dome photometric calibration instrument. **Cite a CBP paper.** The CBP provides a tunable monochromatic calibration source installed inside the Rubin Observatory dome. The optical design is that of a Schmidt reflector telescope run in reverse, with a fiber-fed tunable laser that covers the LSST bandpass from 300nm - 1100nm. The CBP produces spatially uniform projections across all LSSTCam CCDs through projecting light from the CBP onto different positions of the Simonyi Survey Telescope primary mirror. Because the CBP beam only partially illuminates the Simonyi Survey Telescope’s 8.4-meter primary mirror aperture at any single pointing, a sequence of spatially distinct pointings is executed to sample the full pupil. This targeted partial-pupil approach effectively bypasses the atmosphere and eliminates the severe stray

light and ghosting artifacts associated with traditional diffuse flat-screen illumination.

The optical assembly is mounted on an alt-azimuth mechanism on a dedicated platform installed approximately 10 meters above the Rubin Observatory dome floor, which allows the CBP to pivot explicitly about its exit pupil to ensure that beams projected at varying angles map along identical paths into the telescope pupil. The laser beam is fiber-coupled into an integrating sphere equipped with an interior coating to ensure absolute Lambertian diffuse uniformity and eliminate optical memory. At the primary exit port of the integrating sphere (coincident with the CBP telescope’s focal plane), a motorized mask wheel supports five interchangeable masks—such as a single-spot-per-CCD pin-hole configuration—to generate precise illumination patterns across the camera. Absolute flux monitoring is performed at another port of the integrating sphere using a NIST-calibrated silicon photodiode, read out via an electrometer to determine the exact delivered photon dose. **Could cut for brevity - not sure if we want the nitty-gritty instrumental information here**

The data products generated by the CBP are integrated into Rubin DM via specialized processing workflows to enable precise throughput characterization, and contributions from individual optical components (telescope response, filter throughput, and CCD relative quantum efficiency). By cross-referencing the absolute incident flux recorded by the photodiode with the signal detected across the LSSTCam focal plane, the system can characterize the telescope’s response. The data products are highly effective at mapping relative variations across the focal plane, yielding relative quantum efficiency curves that have revealed percent-level systematic offsets from factory-tabulated sensor data. **mention this as an input to FGCM here?** Furthermore, by taking the ratio of telescope signals captured with and without a filter in the optical path, the CBP provides a self-contained measurement of the LSSTCam filter transmission profiles. This is critical for characterizing the chromatic shift of the filter bandpass edges, which vary across the focal plane by several nanometers due to changing angles of incidence across the LSSTCam filters. These data products also allow for the assessment of optical ghosts, electronic crosstalk, and infrared fringing patterns in the z and y bands, while guaranteeing robust calibration coverage during daytime or overcast periods when on-sky observations are not possible.

3.2.4. 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 permanently installed at 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 high-power LEDs selected to span the complete *ugrizy* filter passbands. 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 hardware is housed in a dedicated 590 kg environmentally controlled enclosure on the dome platform to minimize transmission losses across its 20-meter optical fiber runs. Both illumination sources feed into the projector optics to emit an $\sim f/4$ beam at a working distance of 3.2 meters, filling a 74.1 cm diameter monolithic aluminum aspheric reflector mounted via an adjustable hexapod at the top end of the telescope to cast a flat pupil image onto the screen. The output beam’s intensity and spectral content are continuously tracked by an in-dome monitoring suite consisting of a NIST-calibrated silicon photodiode read out by an electrometer alongside two fiber-fed spectrographs covering a combined range of 315-1150 nm.

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. Broad-

band flatfield exposures are acquired every night **Check this, but should be right?** to generate flat frames used during ISR. 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. Conversely, monochromatic flatfield data products from the tunable laser are utilized to map the system’s chromatic throughput as a function of focal plane position. By isolating wavelength-dependent responses of the filters, mirrors, corrector lenses, and sensor quantum efficiency, these monochromatic profiles characterize fine spatial variations across the LSSTCam filters.

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 12), is described in §5.

4.1. Curated Calibrations

4.2. Auxiliary Data Products

4.3. Crosstalk

Describe how the crosstalk matrices are constructed.

4.4. 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.

4.5. Non-Linearity

Non-linearity (NL) describes a deviation from simple proportionality when mapping amplified voltage in-




Figure 4.

egrated 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 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 6 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 metadata. 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 break is the “linearity turnoff” (as indicated in Figure 7). Points above the turnoff are excluded from the NL fit because there is a maximum signal above which the measured counts can no longer be trusted due to hardware limitations at the output of the detector (Figure 3).




Figure 5. Calibration production dependency graph. Arrows indicate data or dependency flow (“x → 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.

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




Figure 6. 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.




Figure 7. Linearity turnoff for a representative amplifier. The turnoff marks the highest signal level below which the exposure-time-to-signal relation remains approximately linear.

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 form 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:  $S_{\text{abs}}(\mu), O, \alpha, \beta \leftarrow \text{MINIMIZE}(\text{Equation 2}, D_i^{\text{abs}}, \mu_i^{\text{abs}}, k_j)$   $\triangleright$  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 8. 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 8. 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 9 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 9. 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):

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

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 10 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




Figure 10. 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.

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 re-

quired 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.

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.

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 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.

**Figure 11.** 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.

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 11 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 ([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.

4.9. Bias

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

4.10. Dark

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

4.11. Flats

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

5. INSTRUMENT SIGNATURE REMOVAL FOR DP2 SCIENCE DATA

Figure 12 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 12 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 without changing their physical meaning.
2. **Dither counts.** A uniform random offset in $[-0.5, 0.5]$ ADU is added to every pixel. This 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




Figure 12. 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.

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 non-linearities 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 readout 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 §7.1.

6. DISCUSSION

Describe how calibrations 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

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”.

7.1. 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 which calibration dataset was used.

7.2. 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.

7.3. 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.

7.4. Certification

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

7.5. Data Monitoring

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

8. LIMITATIONS & KNOWN ISSUES

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

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 12). 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 12, §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: *algo-*

rithms (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 transitions 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, 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, 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, ., 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)