Draft version October 21, 2026
Typeset using LaTeX **twocolumn** style in AASTeX7

https://doi.org/10.71929/rubin/3377440

The Vera C. Rubin Observatory Data Preview 2

Vera C. Rubin Observatory Team,[1] Eric C. Bellm,[2] Pedro H. Bernardinelli,[3] Robert D. Blum,[1] Dominique Boutigny,[4] Massimo Brescia,[5] Alexander Broughton,[6] Audrey Budlong,[7] Rodolfo Canestrari,[8] Neven Caplar,[2] Colin Orion Chandler,[9,2,10] Johann Cohen-Tanugi,[11] Franco Colleoni,[12] Julio Eduardo Constanzo Córdova,[12] Cyrille Doux,[13] Alex Drlica-Wagner,[14,15,16,17] Siegfried Eggl,[18] Anthony Englert,[19] Kevin Fanning,[20] Shenming Fu,[6] Ranpal K. Gill,[12] Enrico Giro,[21] Robinson Godoy,[12] Ivan Gonzalez,[12] Melissa L. Graham,[2,3] Leanne P. Guy,[12] Matthew J. Holman,[22] Minhee Hyun,[23] David Jiménez Mejías,[12] Mario Jurić,[3] Yijung Kang,[6,12] Arun Kannawadi,[24,25] Kshitija Kelkar,[12] Pierre-François Léget,[25] Shuang Liang,[20] Kian-Tat Lim,[20] Huan Lin,[14] Lauren A. MacArthur,[25] Sean Patrick MacBride,[26] Guido Maulen,[12] Samuel Mesquita,[13] James R. Mullaney,[27] Kate Napier,[6] William O'Mullane,[12] Ian Órdenes,[12] Aashay Pai,[15,16] Karla Peña Ramírez,[12] Andrés A. Plazas Malagón,[20,6] Meredith L. Rawls,[2,3] Kevin A. Reil,[20] Gabriele Rodeghiero,[28] Luca Rosignoli,[29,28] David Sanmartim,[12] Cristián Silva,[12] Brian Stalder,[1] Ian S. Sullivan,[2] Diego Tapia,[12] Alessio Taranto,[28,29] Dan S. Taranu,[25] Yousuke Utsumi,[30] and Tianqing Zhang[20]

[1] *NSF-DOE Vera C. Rubin Observatory / NSF NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, USA*
[2] *University of Washington, Dept. of Astronomy, Box 351580, Seattle, WA 98195, USA*
[3] *Institute for Data-intensive Research in Astrophysics and Cosmology, University of Washington, 3910 15th Avenue NE, Seattle, WA 98195, USA*
[4] *Université Savoie Mont-Blanc, CNRS/IN2P3, LAPP, 9 Chemin de Bellevue, F-74940 Annecy-le-Vieux, France*
[5] *Department of Physics "E. Pancini", University Federico II of Napoli, Via Cintia, 80126 Napoli, Italy*
[6] *Kavli Institute for Particle Astrophysics and Cosmology, SLAC National Accelerator Laboratory, 2575 Sand Hill Rd., Menlo Park, CA 94025, USA*
[7] *University of Washington, Dept. of Physics, Box 351580, Seattle, WA 98195, USA*
[8] *INAF Istituto di Astrofisica Spaziale e Fisica Cosmica di Palermo, Via Ugo la Malfa 153, 90146, Palermo, Italy*
[9] *LSST Interdisciplinary Network for Collaboration and Computing, Tucson, USA*
[10] *Department of Astronomy and Planetary Science, Northern Arizona University, P.O. Box 6010, Flagstaff, AZ 86011, USA*
[11] *LPCA, Université Clermont-Auvergne, CNRS/IN2P3, Clermont-Ferrand, France*
[12] *NSF-DOE Vera C. Rubin Observatory / NSF NOIRLab, Casilla 603, La Serena, Chile*
[13] *Université Grenoble Alpes, CNRS/IN2P3, LPSC, 53 avenue des Martyrs, F-38026 Grenoble, France*
[14] *Fermi National Accelerator Laboratory, P. O. Box 500, Batavia, IL 60510, USA*
[15] *Department of Astronomy and Astrophysics, University of Chicago, 5640 South Ellis Avenue, Chicago, IL 60637, USA*
[16] *Kavli Institute for Cosmological Physics, University of Chicago, 5640 South Ellis Avenue, Chicago, IL 60637, USA*
[17] *NSF-Simons AI Institute for the Sky (SkAI), 172 E. Chestnut St., Chicago, IL 60611, USA*
[18] *University of Illinois, Department of Astronomy, 1110 W. Green St., Urbana, IL 61801, USA*
[19] *Department of Physics, Brown University, 182 Hope Street, Providence, RI 02912, USA*
[20] *SLAC National Accelerator Laboratory, 2575 Sand Hill Rd., Menlo Park, CA 94025, USA*
[21] *INAF Osservatorio Astronomico di Trieste, Via Giovan Battista Tiepolo 11, 34143, Trieste, Italy*
[22] *Center for Astrophysics, Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138*
[23] *Stanford University, 450 Jane Stanford Way, Stanford, CA 94305, USA*
[24] *Department of Physics, Duke University, Durham, NC 27708, USA*
[25] *Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA*
[26] *Physik-Institut, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland*
[27] *Astrophysics Research Cluster, School of Mathematical and Physical Sciences, University of Sheffield, Sheffield, S3 7RH, United Kingdom*
[28] *INAF Osservatorio di Astrofisica e Scienza dello Spazio Bologna, Via P. Gobetti 93/3, 40129, Bologna, Italy*
[29] *Department of Physics and Astronomy (DIFA), University of Bologna, Via P. Gobetti 93/2, 40129, Bologna, Italy*
[30] *National Astronomical Observatory of Japan, Chile Observatory, Los Abedules 3085, Vitacura, Santiago, Chile*

(Dated: October 21, 2026)

Corresponding author: Leanne P. Guy, James Mullaney, Alex Broughton

ABSTRACT

We present an early version of Rubin Data Preview 2 (DP2), the second data preview from the National Science Foundation–Department of Energy Vera C. Rubin Observatory. Early Data Preview 2 (EDP2) comprises coadds, detection catalogs, and ancillary data products; eventually, the full DP2 will also include single-epoch images and difference images, but any analysis begun on EDP2 will require no modification. DP2 is derived from observations acquired by the LSST Science Camera (LSSTCam) on the Simonyi Survey Telescope at the Summit Facility on Cerro Pachón, Chile, primarily during the on-sky commissioning campaign between 2025-04-25 and 2025-09-21, supplemented by observations taken between 2025-10-25 and 2026-01-06 that overlap the commissioning footprint. The DP2 footprint comprises the Science Validation wide-area survey, five Deep Drilling Fields, and a number of targeted small-field regions, including Trifid and Lagoon, Prawn, M49, and New Horizons, all observed as part of the Rubin First Look campaign. Each field was imaged in up to six broad photometric bands, *ugrizy*, and coadded to produce deep imaging covering an estimated 3,000 $\deg^2$. The addition of single-visit-only areas expands the total DP2 footprint to an estimated 15,000 $\deg^2$, with coverage in at least one filter. The median per-visit PSF FWHM across the wide-area survey ranges from 1.17 arcsec in the z band to 1.26 arcsec in g and r bands. COSMOS, the deepest field, reaches estimated coadded 5σ depths of u=26 mag, g=26.8 mag, r=26.3 mag, i=26.1 mag, z=25.3 mag, y=23.9 mag. DP2 contains 790 million distinct astrophysical objects, of which 299,343 solar system objects. DP2 is approximately 59.15 TB in size and is available to Vera C. Rubin Observatory data rights holders via the Rubin Science Platform, a cloud-based environment for the analysis of petascale astronomical data. Based on a roughly five-month primary observing baseline and covering only part of the eventual LSST footprint, DP2's area, depth, and multiband coverage nonetheless support a broad range of early science investigations ahead of LSST Data Release 1.

Keywords: Rubin Observatory - LSST

1. INTRODUCTION

The National Science Foundation (NSF)–Department of Energy (DOE) Vera C. Rubin Observatory is a ground-based, wide-field optical-near-infrared facility on Cerro Pachón in northern Chile, built to carry out the Legacy Survey of Space and Time (LSST) (Ivezić et al. 2019a). The observatory is named in honor of Vera C. Rubin, whose observations in the latter half of the 20th century established the first compelling empirical case for dark matter (Rubin & Ford 1970; Rubin et al. 1980). Its 10-yr survey will image the entire visible southern sky every 3–4 nights, reaching a coadded depth of $\sim$ 27 mag in the r-band across an 18,000 $\deg^2$ footprint (Bianco et al. 2022).

At the heart of the facility is the 3.2 gigapixels-gigapixel LSST Science Camera (LSSTCam), mounted on the 8.4 m (6.5 m effective aperture) Simonyi Survey Telescope (Stalder et al. 2024). Together they deliver an instantaneous field of view of 9.6 $\deg^2$ in six broadband filters, *ugrizy* (320–1050 nm), with seeing-limited image quality. The system's étendue of 319 $\deg^2\,m^2$ exceeds that of any previous optical facility by over an order of magnitude, enabling the rapid, deep, all-sky cadence that underpins the LSST's four core science themes: the nature of dark energy and dark matter; a census of the solar system; the transient and variable optical sky; and the structure of the Milky Way (Ivezić et al. 2019a). An automated Data Management System processes tens of terabytes of images per night to produce science-ready data products within minutes of each observation for a global community of scientists.

This paper presents Rubin's third Data Preview, Data Preview 2 (DP2) (NSF-DOE Vera C. Rubin Observatory 2025), and the first to be based entirely on LSSTCam data. It marks the conclusion of the Rubin commissioning program and the transition to full survey operations. DP2 is a reprocessing of all science-grade exposures acquired with LSSTCam during on-sky commissioning and the initial phase of early operations, prior to the formal start of the LSST on 2026-06-29. As such, it represents a natural boundary and provides a comprehensive dataset that captures the on-sky performance of the as-built system at the end of the Rubin construction project.

Two earlier Data Previews preceded DP2 as part of the Rubin Early Science Program (Guy et al. 2026), established to provide the community with early access to data products and services needed to produce high-impact science during the period spanning commissioning through the first Data Release. Data Preview 0 (DP0) was based on simulated LSST-like data (LSST

Dark Energy Science Collaboration (LSST DESC) et al. 2021) and served primarily to stand up early versions of the Rubin Science Platform (RSP) and prepare the community for data access at scale. Data Preview 1 (DP1) (NSF-DOE Vera C. Rubin Observatory 2025) delivered the first data products derived from on-sky observations using science-grade exposures acquired with Rubin Commissioning Camera (LSSTComCam) (Stalder et al. 2022, 2020; Howard et al. 2018; SLAC National Accelerator Laboratory & NSF-DOE Vera C. Rubin Observatory 2024); a 144-megapixel, single-raft instrument used during early commissioning prior to the installation of LSSTCam on the Simonyi Survey Telescope. DP2 supersedes both in scope, fidelity, and scientific depth.

DP2 contains 29,756 science-grade exposures acquired over 158 nights between 2025-04-16 and 2025-09-21, covering a total sky area of 15,000 $\deg^2$ across 85 fields. The data products include raw and calibrated single-epoch images, coadded images, difference images, source detection catalogs, and a range of derived data products. The release totals approximately 59.15 TB and contains around 790 million distinct astronomical objects. Full release documentation is available at https://dp2.lsst.io.

Because DP2 is drawn from a commissioning and early-operations data rather than from the more uniform and systematic LSST, its images and derived products reflect a system that was still being commissioned and tuned. LSSTCam's greater scale and complexity relative to LSSTComCam introduced systematics with no direct counterpart in DP1. For much of the campaign the Active Optics System (AOS) closed loop converged on a slightly mis-calibrated focus target, imprinting a spatially varying defocus across the focal plane that was only corrected on sky in December 2025 and that remains visible in the point-spread function (PSF) of a large fraction of DP2 exposures. Delivered image quality and PSF ellipticity did not consistently meet LSST design requirements in wide-area survey-mode observations by the close of commissioning; and stray- and scattered-light artifacts, sky-background systematics in crowded fields, and a chromatic PSF dependence[31] were identified and only partially mitigated in DP2 processing. These effects, their diagnosis, and the mitigations applied are described in §2 and §6.

Unlike the systematic, rolling revisit pattern that will characterize the LSST, DP2's observing cadence is opportunistic: fields and filters were prioritized to serve AOS and Science Pipelines commissioning rather than to build a representative time series. The Science Validation survey footprint was repeatedly re-scoped over the course of the science validation campaign due to the competing demands of engineering tests and filter availability, which was constrained by hardware issues for part of the campaign. Visit counts and time sampling therefore vary substantially and idiosyncratically from field to field, and DP2 should not be read as a preview of LSST's eventual survey cadence, revisit statistics, nor depth-accumulation pattern.

DP2 was delivered in two phases across the Jul–Dec 2026 window. The first phase, Early DP2 (EDP2), comprised the deep coadd images and all associated catalogs, including new cell-based coadds and the first releases of the ShearObject (§5.3.4, §4.2.8) shape-measurement catalog and the Isolated Star Stellar Motions catalog (§4.2.7). EDP2 was released on 2026-07-27. The second phase, DP2, added the visit-level single-epoch and difference images, completing the full release, and is expected in the Oct–Dec 2026 range. Data products delivered as EDP2 were incorporated into DP2 unchanged, without additional reprocessing. This phased delivery follows directly from Rubin's hybrid cloud architecture (O'Mullane et al. 2024a), in which science data products are stored at the United States Data Facility (USDF) and served to users via the US Data Access Center (DAC) in Google Cloud. Deep coadd images and catalogs are hosted directly in Google Cloud and so do not depend on hybrid-cloud readiness, which allowed them to be released as EDP2. Visit-level images, however, must be streamed from the USDF, and the capacity to do so was not fully in place at the time of the EDP2 release, so they were released later as part of full DP2.

DP2, like all Rubin Data Previews and Data Releases, is subject to a 2 yr proprietary period; during this time access is restricted to Rubin data rights holders (Blum & the Rubin Operations Team 2020) via the RSP (Jurić et al. 2019), Rubin's cloud-based data science platform (Jurić et al. 2019). After the proprietary period, DP2 will be made public, though access through Rubin DACs in the US and Chile will remain limited to data rights holders (Blum & the Rubin Operations Team 2020). Alternative public access mechanisms have not yet been finalized.

The structure of this paper is as follows. Section 2 describes the observatory system and its status during commissioning, and section 3 describes the fields included in DP2, and the observing strategy employed. Section 4 summarizes the data products in the release. The processing pipelines are described in Section 5, followed by data validation and performance characterization in Section 6. Section 7 describes the RSP. Section 8 outlines Rubin's community support model. A summary

[31] LSSTCam has no atmospheric dispersion corrector

and outlook toward the start of formal LSST Data Releases is given in Section 9.

All magnitudes are quoted in the AB system (Oke & Gunn 1983) unless otherwise specified.

2. COMMISSIONING SYSTEM CONFIGURATION AND STATUS

This section describes the components of the Rubin Observatory system and their integration status during the period in which the DP2 data were acquired. Throughout this period the observatory remained under construction: several key components were incomplete, and the performance of the system as a whole was still being characterized. The state of the Simonyi Survey Telescope, LSSTCam and its filter complement, and the calibration system is described first, followed by the data acquisition and processing infrastructure and the delivered image quality. Dedicated engineering campaigns were carried out to test individual system components and to verify their integration into the full system. These campaigns continued as needed throughout the dedicated Science Validation (SV) campaign period described in §3.

2.1. *Simonyi Survey Telescope*

The Simonyi Survey Telescope (Stalder et al. 2024) features a compact three-mirror design incorporating an 8.4-meter primary–tertiary mirror (Single piece of glass for Primary Mirror/Tertiary Mirror (M1M3)) fabricated from a single borosilicate monolith, and a 3.5 m secondary mirror (M2). This compact configuration supports a wide 3.5°field of view while enabling exceptional pointing stability, with slew-and-settle times under 5 seconds.

To achieve the scientific goals of the 10-yr LSST, the observatory must maintain high image quality across the full field of view (Ivezić et al. 2019b). This is accomplished through the AOS (Xin et al. 2015; Megias Homar et al. 2024), which corrects, between successive exposures, wavefront distortions caused by optical misalignments and mirror surface deformations, primarily driven by gravitational and thermal loads.

The AOS comprises open-loop and closed-loop components. The open-loop component corrects for predictable distortions using look-up tables as a function of telescope elevation and camera rotator angle. The closed-loop component addresses unpredictable or slowly varying aberrations using curvature wavefront sensing (Thomas et al. 2023) on the eight corner wavefront sensors of the LSSTCam focal plane, which acquire pairs of intra- and extra-focal "donut" images to measure wavefront errors. Corrections are applied to the M2 and camera hexapods and to the shape of M1M3 via its 156-actuator support system, driving the system towards the target optical state within two visits of each measurement. This represents a significant operational advance over the LSSTComCam campaign, during which dedicated wavefront sensors were unavailable and wavefront errors were instead estimated by defocusing the telescope ± 1.5 mm on either side of focus (NSF-DOE Vera C. Rubin Observatory 2025).

During the LSSTCam commissioning campaign, AOS operations were interspersed with dedicated engineering periods focused on wavefront sensing and optical state characterisation, enabling iterative improvements across all components of the system throughout the campaign (Bechtol 2025). The bulk of the DP2 observations were acquired during the commissioning-era SV surveys, a period during which the AOS was undergoing intensive characterisation and iterative improvement. During this time the system contribution to the delivered Point Spread Function (PSF) Full Width at Half-Maximum (FWHM) still exceeded the design specification (Claver & The LSST Systems Engineering Integrated Project Team 2017; NSF-DOE Vera C. Rubin Observatory 2026a,b). As a consequence, the majority of DP2 observations were taken with a larger system contribution to the PSF size budget than the system is capable of delivering under comparable atmospheric conditions, producing larger PSF FWHMs and shallower point-source depths. The residual optical aberrations also manifest as spatially coherent structure in PSF shape residuals across the focal plane, with implications for weak lensing analyses and PSF model validation. The quantitative impact on delivered image quality is characterised in §3.4, and the AOS improvement program and current performance are described in NSF-DOE Vera C. Rubin Observatory (2026a) and NSF-DOE Vera C. Rubin Observatory (2026b).

2.2. *LSSTCam*

LSSTCam (Corbett et al. 2024; SLAC National Accelerator Laboratory & NSF-DOE Vera C. Rubin Observatory 2025) is a 3.2-gigapixel mosaic camera with a 9.6 deg^2 instantaneous field of view. Its focal plane comprises 25 modular rafts arranged in a 5×5 grid: 21 science rafts, each containing nine 4K×4K Charge-Coupled Device (CCD) sensors, and 4 corner rafts each housing two wavefront sensors and two guider sensors, for a total of 205 sensors: 189 science, 8 wavefront, and 8 guider.

LSSTCam employs sensors from two vendors: Teledyne-e2v semiconductor company (E2V) sensors populate the central science rafts, and Imaging Technology Laboratory (UA) (ITL) sensors populate the pe-

ripheral science rafts. Each sensor is subdivided into 16 amplifiers, each reading 512×2048 pixels in parallel. Figure 1 shows the LSSTCam focal plane layout, viewed from above looking down through the camera lenses, including the placement and numbering scheme of sensors ("S") and amplifiers ("C"); a full description is given in Plazas Malagón et al. (2025); Plazas Malagón et al. (2026).

During the commissioning campaign, only 188 of the 189 science sensors were operational; sensor R30/S12 was non-operational throughout and is excluded from all DP2 data products. The eight heavily vignetted corner science sensors (detectors), R01/S00 (0), R03/S02 (20), R10/S00 (27), R14/S02 (65), R30/S20 (123), R34/S22 (161), R41/S20 (168), and R43/S22 (188), are also excluded from all DP2 processing, even though they were functional. Science sensors (detectors) R03/S20 (24), R30/S10 (120), R30/S11 (121), and R42/S00 (172) are additionally excluded from data acquired before 2025-07-03 as the calibration data required to process them were not yet available. Any additional exclusions that were applied during the processing stage are described in §5.

The median amplifier loses approximately 0.45% of its area to static defects, with the mean across amplifiers closer to 0.7% (Bechtol 2025). Overall, more than 99% of the active focal-plane area within the calibrated field is usable (Bechtol 2025). Optical vignetting limits the fully calibrated field radius to 317 mm; a partially vignetted zone extends to 350 mm, beyond which pixels are excluded from science products.

No new static defect classes were identified on either ITL or E2V sensors during commissioning beyond those already identified during commissioning with LSSTComCam (Vera C. Rubin Observatory Team 2026) or in the laboratory (LSST Camera Team et al. 2025; Bechtol 2025). A subset of ITL sensors exhibits "vampire" pixel anomalies and locally elevated readout noise consistent with LSSTComCam commissioning. Dynamic masking, including treatment of E2V edge bleeds and occasional bright-star flare artifacts at extreme detector edges, is ongoing (Bechtol 2025).

2.2.1. *Filter Complement*

LSSTCam supports imaging with the same six broadband *ugrizy* filters, spanning 320–1050 nm, as LSSTComCam, using an identical filter design. Its filter exchanger can hold five filters at a time, as opposed to only three for LSSTComCam.

Unlike the LSSTComCam bandpasses, which were built entirely from fiducial mirror, lens, filter, and detector quantum-efficiency curves in the absence of any device-specific measurements, the DP2 standard bandpasses incorporate real measurements where available. The fiducial mirror and lens throughputs are still used, since detector- and mirror-level measurements have not yet been incorporated into DP2. The filter throughput, however, is now the focal-plane average of vendor filter scans, using the same underlying filter design as LSSTComCam but capturing real variation across the field of view. The detector quantum efficiency is likewise the focal-plane average over all the ITL and E2V detectors combined.

ITL and E2V sensors are not identical: ITL devices have somewhat higher quantum efficiency in the red, so the focal-plane-averaged LSSTCam response is marginally lower than the all-ITL response used for LSSTComCam. This chromatic difference is most apparent in the g band, visible even by eye as a two-tone effect across the focal plane. u-band throughput is dominated by the atmosphere and filter rather than the detector, and the redder bands are comparatively insensitive to the vendor difference. In absolute terms the ITL/E2V difference is small, much smaller than, for example, the difference between LSSTCam and other instruments such as Dark Energy Camera (DECam), since the two vendors' filters are essentially the same design. It is nonetheless large enough to matter for millimagnitude-level photometric calibration.

Individual per-detector quantum-efficiency curves have not yet been released. The standard passband distributed with DP2 is therefore the focal-plane average, consistent with how the survey is calibrated and appropriate for use with the deep coadds. The full-system throughput of the six LSSTCam filters is shown in Figure 2 and the corresponding transmission curves are provided as a DP2 data product(§4.4.3).

2.2.2. *Stray Light*

The compact, 3-mirror design of the Simonyi Survey Telescope (§2.1) is especially sensitive to stray and scattered light[32]. The most pronounced and extended artifacts appear in LSSTCam images when the telescope's boresight is near bright objects such as the Moon, bright stars, planets, or unintended light sources on the sky. Stray light contaminates all astronomical imaging to some level, but LSST's low-surface-brightness science is especially sensitive to it.

[32] Light that does not follow the nominal optical path within the telescope and instrument, originating from scattering off optical and non-optical surfaces or from ghost paths in refractive components such as lenses and filters

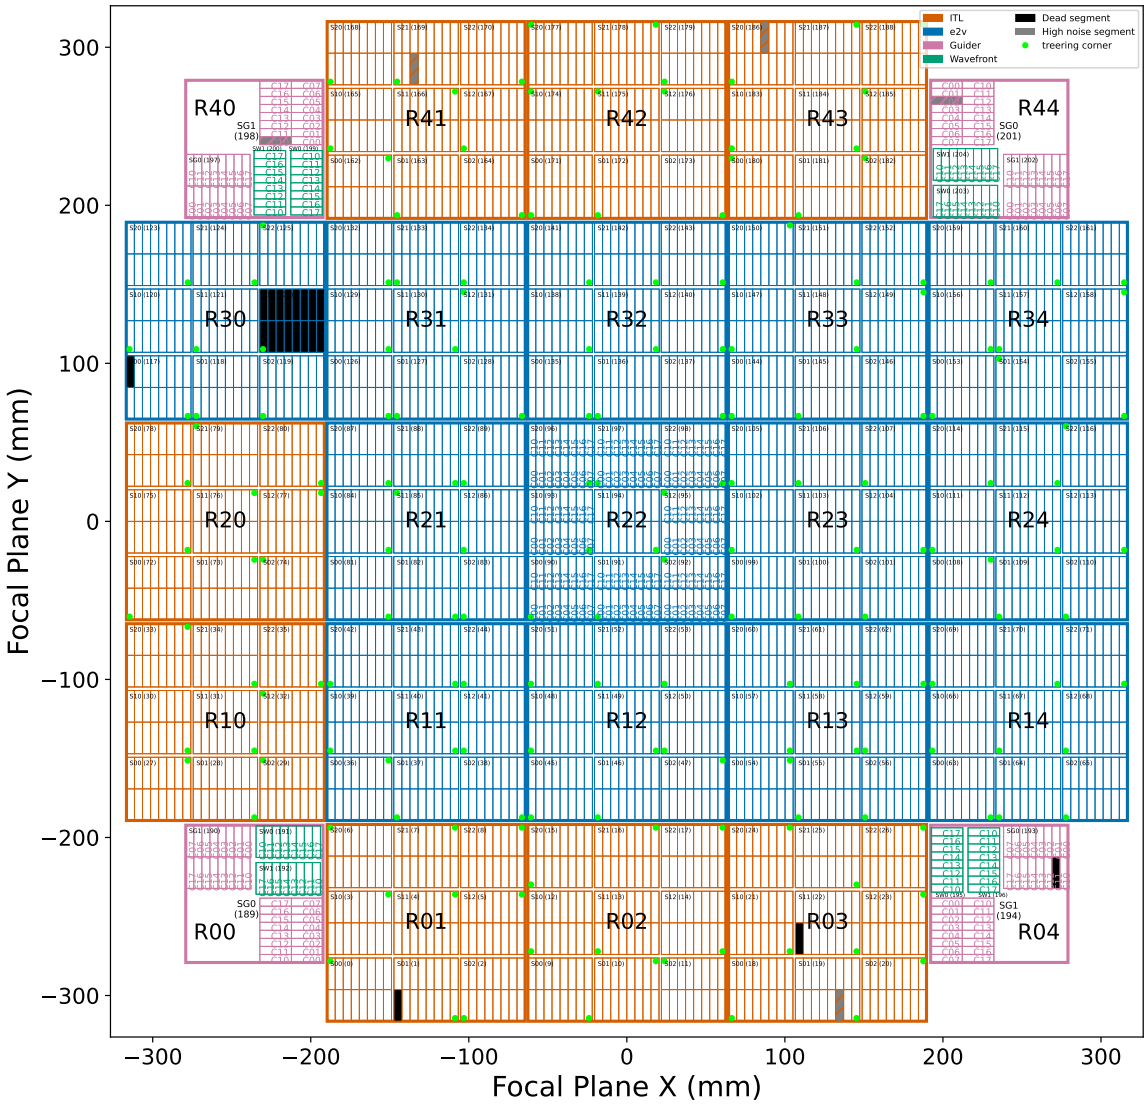


Figure 1. LSSTCam focal plane layout, showing the placement and numbering scheme of sensors ("S") and amplifiers ("C"). The view is from above, looking through the LSSTCam lenses onto the focal plane. Each sensor contains 16 amplifiers. Each raft is composed of a 3×3 array of sensors. The detector number for each sensor is indicated in parentheses.

A significant fraction of these artifacts traced back to the Observatory dome, which lacked its Light-Wind Screen (LWS) to restrict the acceptance angle along the telescopes's elevation axis; many artifacts are expected to resolve once the LWS is comissionined, whereas others will require additional baffling. Some artifacts detected were expected based on an initial stray light analysis in FRED® (Ellis 2009), which correctly predicted direct paths from the primary (M1) and tertiary (M3) mirrors that bypass the secondary mirror and hit the LSSTCam optics and focal plane directly. Other artifacts were new and unexpected, requiring visual inspection, ray tracing simulation, and on-sky tests to confirm their origins (Rodeghiero et al. 2026).

An initial census of stray light artifacts observed in LSSTCam images during initial commissioning (April 2025–May 2026) is reproduced in Table 1 from (Rodeghiero et al. 2026). About 70% of features come from identified, confirmed sources, and 80% already have a possible mitigation strategy. Occurrence rates are based on visual classification of features visible by eye; for some, fractional area is a rough estimate due to variation across frames.

The most frequent and impactful feature, the *scratched tape* feature (Drlica-Wagner et al. 2026), originates from the M1 direct path that skips M2 and M3 reflections, and was visible in ∼7.4% of image frames with an off-axis range of visibility between ∼19.°8 and 22°. An annular extension of the telescope's mid-level

Table 1. Occurrence rate of stray and scattered light artifacts identifiable by visual inspection of ∼116,000 full focal plane mosaic images. Rates are not expected to be complete nor pure; impacted area uses average feature size. The most prevalent artifact, the *scratched tape*, was effectively mitigated through installation of the mid-level baffle extension (Drlica-Wagner et al. 2026). * Fractional area of the artifact is variable; † artifact origin to be confirmed; ϒ preliminary estimate.

Artifact	Source	Mitigation	Occurrence Rate	Fractional Area
Scratched Tape	M1 direct path	Mid-Level Baffle Extension	7.39%	2.2% *
Smudge	M1 direct path	Mid-Level Baffle Extension	0.19%	0.28%
Muddy Shoe	M1 direct path	Mid-Level Baffle Extension	0.15%	2%*
Pillow	Unknown	LWS installation	0.26%	9%*
Nibbled Disc	Unknown	Mid-Level Baffle Extension	0.10%	1.06 - 10%*
L2 Bananas	L2 support pads	LWS installation	0.73%	0.7%
Rivet Holes	Holes in Top-end Baffle	Sealing of holes	0.45%	–
Sharp Arcs	L2 edges †	TBD	TBD	1% ϒ
Submarine	L1/L2 ghosts †	LWS installation (TBC)	0.7% ϒ	8.5%*
L3 Horseshoe	L3 chamfer & mount	Additional L3 edge baffle	0.09%	8.8%
All artifacts	—	—	10.00%	—
Post Mid-Level Baffle Extension installation		—	2.6%	—

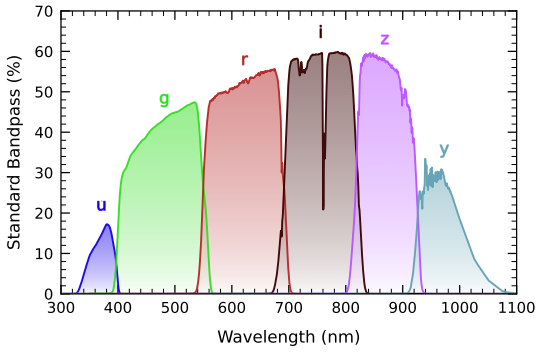


Figure 2. LSSTCam standard bandpasses, illustrating full-shown stem throughput. The bandpasses include a standard atmosphere at an airmass of 1.2, telescope optics, camera surfaces, and mean ITL and E2V detector quantum efficiency. The corresponding transmission curves are provided as a DP2 data product.

baffle ring, installed between February and April 2026, elimintated this feature from LSSTCam images (Drlica-Wagner et al. 2026). Other features such as the *L3 Horseshoe*, arise from edge diffraction at LSSTCam's tertiary lens (L3) near the field-of-view edge (∼ 1.°99 and 2.°23). Ray tracing simulations also revealed scattered light from the Secondary Mirror (M2) optical baffle, whose residual reflectance in the longer wavelength bands (i, z, y) can reach up to 30%.

Despite the preventative measures taken, ∼10% of early image frames suffer from stray light contamination, dropping to ∼2.6% following the mid-level baffle extension installation (Drlica-Wagner et al. 2026). The most severe issues were structured artifacts on the focal plane, largely stemming from the incomplete status of the LWS. Rubin is implementing new physical baffling and re-painting to address the remaining artifacts (Pollek et al. 2026; Taranto et al. 2026).

2.3. *Calibration System*

Unlike the LSSTComCam campaign, during which the Rubin calibration system had not yet been installed and flat fielding relied entirely on twilight flats (Vera C. Rubin Observatory Team 2026), the LSSTCam commissioning campaign had access to the Rubin flat field illumination system (Fagrelius & Rykoff 2025). The baseline calibration plan for the as-built system is described in Fagrelius & Rykoff (2025); a full description of Instrument Signal Removal (ISR) procedures and calibration data products is given in Plazas Malagón et al. (2025).

The calibration system comprises a dome flat field screen, a Collimated Beam Projector (CBP), a National Institute of Standards and Technology (USA) (NIST)-traceable silicon photodiode electrometer, and the AuxTel for atmospheric monitoring. During the commissioning campaign, the CBP was deployed at scale on LSSTCam for the first time, providing in-situ measurements of the transmission profiles of five of the six LSST broadband filters, *ugrizy*, across multiple primary mirror pupil positions, and characterising the angle-of-incidence-dependent blue-shift of the filter bandpass edges across the focal plane (Amouroux et al. 2026).

Achieving absolute throughput calibration at the millimagnitude level required by the LSST requires a more precise characterisation of the CBP transmission and a partial-pupil vignetting model; both are the subject of ongoing work (Amouroux et al. 2026).

Calibration products generated for DP2 processing include bias frames, dark frames, dome flat fields, Photon Transfer Curves (PTCs), cross talk matrices, brighter-fatter kernels, linearity corrections, and Charge Transfer Inefficiency (CTI) models (Plazas Malagón et al. 2025). Dome flat fields were used as the primary flat field source, supplemented by a star-flat-based illumination correction and, where dome illumination was unavailable, twilight flats (Plazas Malagón et al. 2025). Certified calibration products are versioned in the USDF archive and transferred to the summit as needed. The full suite of calibration and throughput data products included in DP2 is described in §4.

2.4. *Data Management System*

The Rubin Data Management Subsystem (DMS) is responsible for the acquisition, transfer, archiving, processing, and distribution of all data produced by the observatory. It operates as an automated end-to-end system with minimal human intervention, ingesting raw pixels from the LSSTCam focal plane within seconds of shutter close and delivering science-ready data products to a global community of researchers within 120 secs of shutter close.

The data products in DP2 are the result of Data Release Processing (DRP) (§5), which reprocesses the full accumulated image archive to produce the deep, well-calibrated data products that form the basis of each Data Release, and are delivered to the community via the RSP (§7). During the LSSTCam commissioning campaign, the DMS was operated at full scale for the first time, transferring, archiving, and processing all LSSTCam images, and validating the end-to-end data flow from summit to RSP.

2.4.1. *Data Acquisition, Transfer and Archive*

Following each exposure, the Data Acquisition System (DAQ) system at the summit stores the raw image data as a set of 205 files per exposure: 189 science sensors[33], 8 wavefront sensors, and 8 guiding sensors on the LSSTCam focal plane. Uncompressed, a full science exposure totals ~ 8.2 GB; after compression, ~ 4 GB.

At the summit, each exposure is processed in near real time by the Rapid Analysis pipeline, performing ISR (§5.1.1) and source detection (§5.1) on a 30–60 second timescale to provide quick-look image diagnostics and support real-time AOS assessment. The resulting per-visit quality metrics formed the basis of the visit quality characterisations used to define the DP2 coadd inputs. Summary visit and detector-level metadata are available to users as DP2 data products (§4.5) and are also available in an opsim-format SQLite database §7.5.2).

2.4.2. *LSST Science Pipelines*

The DP2 science-ready data products are the result of dedicated DRP campaigns executed at the USDF at SLAC National Accelerator Laboratory (SLAC). DP2 processing was done using version v30.0.8 of the LSST Science Pipelines (Rubin Observatory Science Pipelines Developers et al. 2026) and follows a four-stage structure. Stage 1 performs single-visit processing, including ISR, source detection, astrometric and photometric calibration, and PSF characterisation. Stage 2 performs global recalibration, Gaussian-process modelling of atmospheric turbulence, proper motion and parallax fitting, and PSF re-estimation using a standard star set. Stage 3 performs coaddition and measurement on coadded images, including the generation of the Object catalog. Stage 4 performs Difference Image Analysis (DIA) and solar system object processing, and the generation of light curves via forced photometry. Each stage underwent continuous data quality monitoring and a formal verification and validation review before processing advanced to the subsequent stage. A full description of the processing pipelines, algorithms, and configuration used to produce DP2 is given in §5.

3. SCIENCE PROGRAM OBSERVATIONS

The science program observations acquired during Rubin commissioning, which form the basis of DP2, spanned three distinct phases. The first comprised small-field observations acquired shortly after the LSSTCam First Photon milestone in support of the Rubin First Look event (3.1). The second was the commissioning-era SV survey campaigns, which ran from 2025-04-16 to 2025-09-21 (3.2). The end of the SV campaign was followed by a four-week engineering downtime, after which the observatory was formally handed over from the Construction project to Rubin Operations on 2025-10-25. The third phase comprises pre-LSST Early Operations observations from that date through the DP2 cutoff date of 2026-01-06 (3.3). Early Operations observations continued for approximately six months after the DP2 cutoff, until the start of the LSST on 2026-06-29. All science program observations were scheduled by the Feature-Based Scheduler (FBS), which

[33] Only 188 science sensor files were generated per exposure as sensor R30/S12 was non-operational throughout; see §2.2.

selected visits adaptively to balance sky coverage, filter distribution, airmass, and lunar phase constraints (Jones et al. 2026). Because the program was conducted while the system was still being commissioned and optimized, visit acquisition was driven by engineering and commissioning priorities rather than by a scientifically optimized observing strategy. Engineering observations are not scientifically useful and are not included in DP2.

Figure 3 shows the cumulative distribution of visit MJD for all DP2 visits, combined across all filters. Notable dates are indicated: the DP2 start date, 2025-04-25; the end of the DP2 data collection window, 2026-01-06; a phase of rapid data acquisition (July 1–25, 2025); and the end of the engineering shutdown, when observations resumed at the start of Operations, 2025-10-25. The sky distribution of DP2 visits is shown in Figure 4. The observation cadence during DP2 differs significantly from that of the planned LSST. Rather than following the cadence intended for the full survey, visit acquisition during this period was driven by engineering and commissioning priorities. As a result, the distribution of visits in time reflects the sequence of commissioning activities and engineering interventions rather than a scientifically-optimized observing strategy.

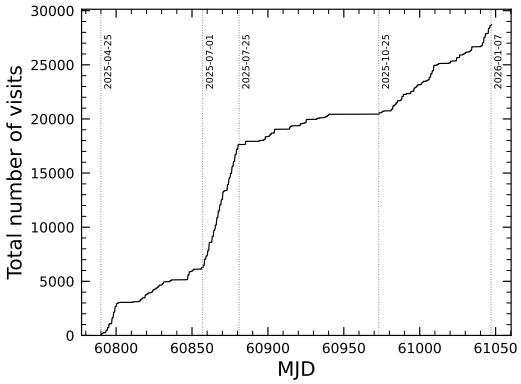


Figure 3. Cumulative distribution of visit MJD for all DP2 visits, combined across all filters. Vertical dotted lines mark notable dates: the DP2 start date, 2025-04-25; the end of the data collection window, 2026-01-06; a phase of rapid data acquisition (July 1–25, 2025); and the end of the engineering shutdown, when observations resumed 2025-10-25

Data collected between the end of the DP2 observing window on 2026-01-06 and the formal start of the LSST on 2026-06-29 will be processed and made available to LSST data rights holders through the same RSP services. The format and timing of that release are still to be determined; planning will begin once DP2 has been delivered (Guy et al. 2026)

Table 2. Median 5σ coadd detection limits per band per small-field science verification region in DP2.

Field Code	Band					
	u	g	r	i	z	y
Abell_2764	···	25.80	···	22.72	···	···
M49	25.57	26.87	26.59	25.95	···	···
Prawn	25.14	26.07	25.82	25.26	23.87	···
Trifid-Lagoon	26.09	26.76	25.96	25.34	23.61	22.34
New_Horizons	25.21	26.29	26.33	25.93	25.14	23.61
Rubin_SV_212_-7	···	26.95	26.72	26.25	···	···
Rubin_SV_216_-17	···	26.32	26.08	25.93	···	···
Rubin_SV_225_-40	26.06	27.09	26.53	26.16	25.34	23.89
Rubin_SV_280_-48	23.79	24.80	24.67	24.67	23.53	···
Rubin_SV_300_-41	···	24.59	···	···	23.65	···
Rubin_SV_320_-15	24.24	25.06	24.90	25.07	24.31	22.68

3.1. *Small-Field Observations*

Several small fields were observed prior to the start of the SV surveys, with some observations overlapping the early portion of the SV survey period. All small-field observations, including those from the First Look campaign, are included in DP2. Observations were typically acquired in sequences of at least 10 visits per bandpass, cycling through the available filters, with field selection guided by commissioning priorities and community input. Several of these fields, including Cosmic Evolution Survey (COSMOS), M49, the Prawn Nebula region, and Rubin_SV_225_-40, received substantial numbers of visits and are among the most deeply observed regions in DP2. The COSMOS field was first observed as part of the small-field campaign but was not accessible during the main SV survey period; additional observations were later acquired during the Early Operations phase.

Table 2 reports the 5σ point–source depths for coadded images per field and per band, where coverage in a band is nonzero.

3.2. *Science Validation Surveys*

The SV surveys began on 2025-04-16 and comprised two interleaved components within a single FBS configuration, a wide survey and a deep survey.

3.2.1. *The Wide Survey*

The Wide Survey initially covered a $\sim 3{,}000\,\mathrm{deg}^2$ footprint tracing the ecliptic plane from the dense regions of the Galactic Bulge through low-dust regions of the planned LSST Wide-Fast-Deep (WFD) footprint. Approximately halfway through the SV period, on 2025-

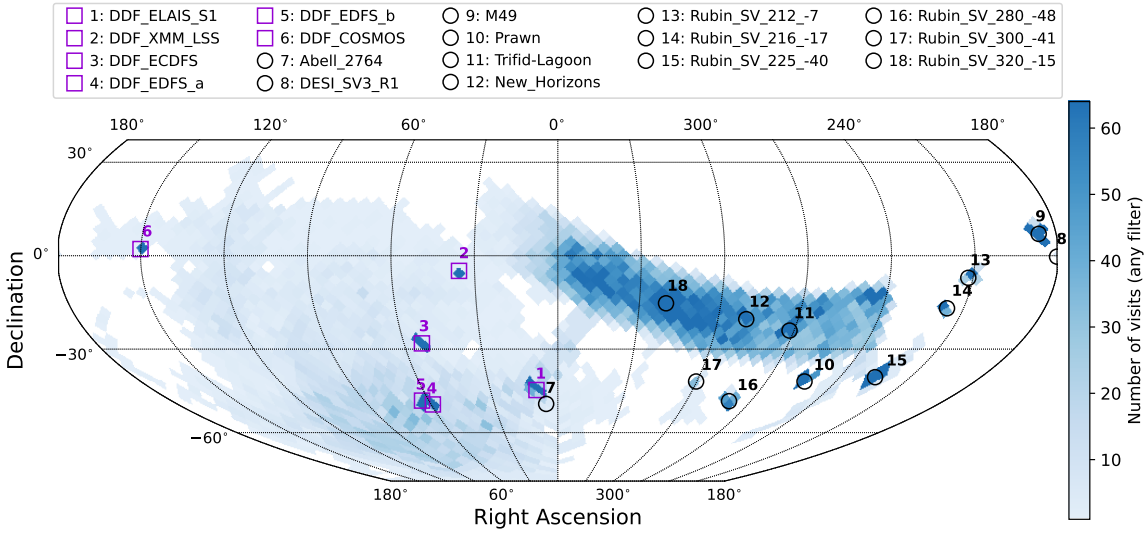


Figure 4. Sky distribution of DP2 visits across the three components of the observing program. Small-field survey areas are marked with black circles and Deep Drilling Fields (DDFs) with violet squares. The SV wide survey covers 3,000 deg^2 tracing the ecliptic plane, with visit density increasing toward the 300 deg^2 core. Pre-LSST Early Operations observations appear as sparser, predominantly *i*-band coverage, extending across a larger area of the LSST footprint.

08-18, the footprint was consolidated to $\sim$750 deg^2 and observed in the *griz* bands only. This reduced footprint is a 10°-wide band centred on the ecliptic, defined by ecliptic longitude $285° < \lambda < 360°$ and ecliptic latitude $|\beta| < 5°$. On 2025-09-20, in the final week of commissioning, the footprint was further reduced to the $\sim$300 deg^2 area covered by newly deployed templates ($300° <$ R.A. $< 324°$, $-26° <$ decl. $< -10°$, and $\beta < 5°$), to concentrate visits and reach the depth required for difference imaging tests. As a result, median visit counts and coadded depth increase substantially toward the inner regions of the footprint.

Table 3 provides a description of the WFD areas observed during the campaign.

Table 4 provides the estimated m_5 coadded depths within each of the subsets of the wide SV area.

3.2.2. *The Deep Survey*

The Deep Survey targeted four of the five planned LSST DDFs: Extended Chandra Deep Field-South Survey (ECDFS), Euclid Deep Field South (EDFS) (fields A and B), ELAISS1, and COSMOS. The FBS interleaved DDF observations with the wide survey, scheduling visits to each field according to the prescribed DDF observing cadence. ELAISS1 received the largest number of visits, 539 across all six *ugrizy* bands, and is the only DDF observed in all bands during the SV survey. COSMOS was not accessible during the SV period as described above. The remaining DDFs were observed primarily in *griz*, as the *u* and *y* filters were unavailable after 2025-08-10 due to a hardware issue with the filter exchanger. The dither pattern for the DDFs evolved during the campaign; early observations used only small translational dithers with no intra-night rotation, while later observations introduced both translational and rotational intra-night dithers.

Table 5 reports the 5σ point–source depths for coadded images per field and per band, where coverage in a band is nonzero.

3.3. *Early Operations Observations*

Following the transition to Operations on 2025-10-25, FBS-driven observations resumed during the Early Operations optimization period. Observations during this period, which ran from 2025-10-25 through 2026-01-06, were a mixture of engineering activities and FBS-driven survey observations aimed at improving image quality and survey speed in preparation for the start of the LSST, and building difference-imaging templates across the full LSST footprint.

The pre-LSST observations are characterised by a lower and less regular cadence than the SV surveys, reflecting frequent engineering interruptions. The FBS was configured with a single-visit template tier as a scheduling priority, and on any given night would typically concentrate observations in a single band rather than interleaving filters, to maximise sky coverage for template building. The *i* band was preferred due to its high throughput and suitability for difference imaging. A concentrated block of *i*-band observations in late October 2025 reflected deliberate AOS engineering tests that held the filter fixed to isolate optical state changes from filter-induced perturbations, producing the prominent *i*-band coverage regions visible in Figure 4. Pre-

Table 3. Definitions of the successive footprints of the SV Wide Survey. λ and β are ecliptic longitude and latitude in the barycentric true ecliptic frame. Areas are computed on a HEALPix grid with $N_{\rm side} = 256$. About 77% of the 300 deg^2 footprint lies within the 750 deg^2 footprint and 99% within the 3,000 deg^2 footprint.

Footprint	Definition	Extent	Area (deg^2)
3,000 deg^2	$\lvert\beta\rvert < 10°$; $\lambda > 240°$ or $\lambda < 40°$	20° wide in β, 160° long in λ	3,184
750 deg^2	$\lvert\beta\rvert < 5°$; $\lambda > 285°$	10° wide in β, 75° long in λ	749
300 deg^2	300° < R.A. < 324°; −26° < decl. < −10°; $\beta < 5°$	$\lambda \approx 297°$–$323°$, $\beta \approx -11°$ to +5°	305

Table 4. Median 5σ coadd detection limits per band in disjoint areas of the wide-area SV survey footprint: the 300 deg^2 core, the 750 deg^2 area outside it, and the 3,000 deg^2 area outside both in DP2

Field Code	Band					
	u	g	r	i	z	y
3,000 deg^2	24.16	25.21	24.73	24.32	23.68	22.47
750 deg^2	24.11	25.32	25.07	24.68	23.85	22.74
300 deg^2	24.27	25.00	24.84	24.75	24.03	22.85

Table 5. Median 5σ coadd detection limits per band per Deep Drilling Field in DP2. The XMM_LSS field does have visit observations, but is not included in the deep coadd area for DP2.

Field Code	Band					
	u	g	r	i	z	y
ELAIS_S1	25.35	26.67	26.04	26.22	25.27	23.81
ECDFS	24.39	25.96	25.67	26.00	24.92	22.51
EDFS_a	23.64	25.75	25.61	25.79	24.71	23.22
EDFS_b	23.71	25.86	25.62	25.80	24.74	23.19
COSMOS	25.87	26.84	26.51	26.09	25.15	23.81

LSST visits not overlapping existing SV survey regions were not combined into coadds but are included in the DP2 visit dataset.

3.4. *Delivered Image Quality*

The delivered image quality at Rubin Observatory reflects contributions from the atmosphere, dome environment, telescope optics, AOS, and LSSTCam. The design specification for the median delivered PSF FWHM, referenced to zenith at 550 nm, is 0.″7 (Ivezić & The LSST Science Collaboration 2018), corresponding to the measured median atmospheric seeing of 0.″72 at Cerro Pachón (NSF-DOE Vera C. Rubin Observatory 2026b) added in quadrature with a total system contribution of 0.″35 (NSF-DOE Vera C. Rubin Observatory 2026b). During the period covered by the DP2 dataset, the system had not yet consistently met this design specification (NSF-DOE Vera C. Rubin Observatory 2026b; Bechtol 2025). The primary excess contributions are from the optical system, including residual AOS errors and mirror thermal gradients, as well as from dome seeing. The current total system contribution to PSF FWHM is approximately 0.″55–0.″60, compared to the design allocation of 0.″40. At the median survey airmass of 1.2, an atmospheric contribution of 0.″72 at zenith corresponds to ∼0.″80; combined in quadrature with the current system contribution, this is consistent with the observed median delivered image quality of ∼1.″0 PSF FWHM in the *riz* bands during the DP2 observations (NSF-DOE Vera C. Rubin Observatory 2026a).

Figure 5 shows the cumulative distribution of the median PSF FWHM (in arcseconds) per visit computed across all detectors, for each band.

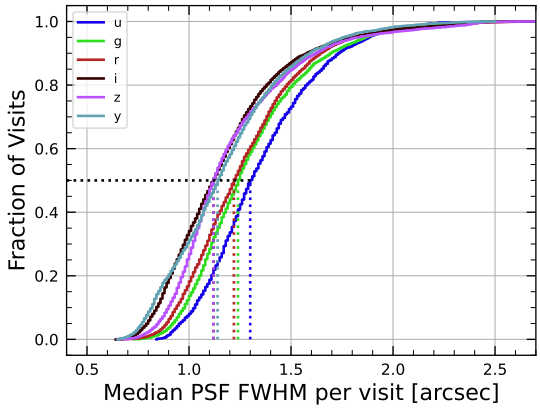


Figure 5. Cumulative distribution of the median per-visit seeing (PSF FWHM) per band. All regions included.

4. OVERVIEW OF THE CONTENTS OF RUBIN DP2

In this section we describe the data products that are included in DP2. Data rights holders can access all of these data products via one or more of the data services outlined in §7. For programmatic reasons, the precise names of the data product can differ slightly depending on which data service is being used to retrieve them. For this reason, we use descriptive names throughout this section and provide their equivalent names used in the data services in Table 6. We note that some data products are not available in all data services; such cases are indicated in Table 6.

As with DP1, all the data products described in this section fall under one of five categories:

- Science Images, including pre-processed Raw Images, Visit Images, and Deep Coadd images;
- Catalogs of astrophysical Sources and Objects detected and measured in the aforementioned science images. We also provide the reference catalog that was constructed to astrometrically and photometrically calibrate the data during processing;
- Maps, which provide non-science-grade visualisations of the data;
- Ancillary Data Products, such as the values of the configuration parameters that were used during data processing, processing logs, and calibration data products;
- Metadata, in the form of tables containing information about each visit and visit image, such as visit pointing coordinates, exposure times, and data quality summary statistics.

We expect Science Images and Catalogs to be the most commonly used categories for science, while Maps, Ancillary Data Products, and Metadata remain important for facilitating science exploitation via visual inspection, as well as for transparency, Education and Public Outreach (EPO), media, and citizen-science activities such as Zooniverse. While the categories of data products in DP2 are the same as in DP1, this second Data Preview contains a wider range of data product types compared to DP1. This is principally due to the LSST science pipelines being in a more mature state than for DP1. Data product types that are new to DP2 are identified as such in their descriptions below.

As with DP1, the LSST science pipelines use a skymap to structure how processing is organized spatially. The skymap divides the sky into a series of equal-area tracts, which are then subdivided into a number of patches. The tracts are arranged in rings of constant decl., starting at the southern pole and having an index number that increases first by increasing R.A. and then by increasing decl.. The skymap that was generated for DP2 processing contains a total of 18,938 tracts spanning the whole sky, 2174 of which were covered during the science validation observing campaign, including the small field surveys. Each DP2 tract covers approximately $2.8\,\mathrm{deg}^2$, and contains 10 $\times$10 patches, both the same as for the DP1 skymap. Tracts are defined to have equal area, and so the amount by which neighboring tracts overlap varies with decl., with tracts near the celestial poles overlapping more than those near the celestial equator. The minimum area of overlap between tracts is $1\farcm0$. Within a tract, the area of overlap between neighboring patches is a constant $60\farcs0$. As a result, each patch covers an area of $0.034\,\mathrm{deg}^2$, which is larger than the tract area divided by the number of patches in a tract. Excluding the overlap region, the area covered uniquely by a single patch is $0.028\,\mathrm{deg}^2$.

Although the tract and patch configuration is unchanged between DP1 and DP2, the two skymaps differ at the cell level for Deep Coadds. Patches are further subdivided into cells for the sole purpose of enabling the construction of cell-based coadds with a continuous PSF (§4.1.3). Unlike patches, cells are not exposed to the user, and one cannot specify nor retrieve an individual "cell" from an image. Instead, cells are used during the construction of what are termed cell-based coadds. A region of an input image is included in a given cell if, and only if, it covers the entire cell. This guarantees that every pixel within a cell is built from exactly the same set of input images, so the PSF is well-defined and effectively constant across the cell, with discontinuities confined to the cell boundaries. Such cell-based coaddition was not implemented during the production of DP1, and was only used to produce the Deep Coadds in DP2. Template Coadds in DP2 do not utilize cells. In the case of DP2 Deep Coadds, there are 22 $\times$22 cells per patch, including the patch overlap region (the patch overlap region is 1 cell wide). There is no overlap between neighboring cells.

4.1. *Science Images*

Science images are images taken with the intent of scientific analysis of astronomical objects, as opposed to calibration images (e.g., bias, darks, flats) or engineering images, which serve to calibrate or test the instrument and telescope regardless of whether they happen to be pointed at the night sky. DP2 provides all images in the acquisition and processing chain, from the raw, unpro-

Table 6. Summary of dataset types released in DP2, together with the names by which they are referred to within the Qserv database and the Data Butler.

Dataset Type	Qserv Name	Butler Name
Raw Images	···	`raw`
Visit Images	···	`visit_image`
Deep Coadd	···	`deep_coadd`
Template Coadd	···	`template_coadd`
Difference Image	···	`difference_image`
Object	dp2.Object	`object`
Isolated Star Stellar Motions	dp2.IsolatedStarStellarMotions	`isolated_star_stellar_motions`
Object Shear	dp2.ShearObject	`object_shear_all`
Source	dp2.Source	`source`
Forced Source	dp2.ForcedSource	`object_forced_source`
Forced Source at DIA Object	dp2.ForcedSourceOnDiaObject	`dia_object_forced_source`
DIA Object	dp2.DiaObject	`dia_object`
DIA Source	dp2.DiaSource	`dia_source`
Solar System Source	dp2.SSSource	`ss_source`
Solar System Object	dp2.SSObject	`ss_object`
MPC Orbit	dp2.mpc_orbits	···
Current Identifications	dp2.current_identifications	···
Numbered Identifications	dp2.numbered_identifications	···
Calibration Reference Catalog	···	`the_monster_20250219`
Survey Property Maps	···	`deepCoadd_<PROPERTY>_consolidated_map_<STATISTIC>`
HiPS Maps	···	···
Skymap	···	`skyMap`
Coadd Patches	dp2.CoaddPatches	···
Standard Passbands	···	`standard_passband`
Object Scarlet Models	···	`object_scarlet_models`
Deep Coadd Input Summary	···	`deep_coadd_input_summary`
Compare Warp Artifact Mask	···	`compare_warp_artifact_mask`
Camera	···	`camera`
FGCM Lookup Table	···	`fgcmLookUpTable`
Pretrained Model Package	···	`pretrainedModelPackage`
Run Provenance	···	`run_provenance`
Calibration datasets	···	`bias`, `dark`, `flat`, `bfk`, `crosstalk`, `cti`, `defects`, `linearizer`, `ptc`
Visit	dp2.Visit	`visit_table`
Visit Detector	dp2.VisitDetector	`visit_detector_table`
Visit Summary	···	`visit_summary`

cessed science and calibration images through to fully calibrated visit images, coadds, and template images. Most science analyses will start from the Deep Coadd images and the single-epoch Visit images. While it is possible for users to recreate their own Processed Visit Images (PVIs) from the raw and calibration data, we do not expect anyone will need to do so. One important difference between DP1 and DP2 is the format of the Science Images (Bosch 2026). In the case of DP2, the format of the datasets are such that they are more compatible with external packages, in particular Astropy. It is, however, straightforward to convert back to "legacy" DP1-style datasets for backwards-compatability.

4.1.1. *Raw Images*

Raw Images are produced by the camera following an exposure. They are the initial input science images to the pipelines and have not undergone any processing. Each active detector in LSSTCam (§2.2) produces one Raw Image per exposure. Each exposure therefore results in up to 180 Raw Images, although some exposures taken during the observing campaign resulted in fewer than 180 Raw Images due to the problematic detectors described in §2 and occasional readout failures.

All LSSTCam detectors have 16 amplifier segments, arranged in two rows of eight, each reading out its own ∼1 megapixel region in parallel to enable the fast (∼2 second) full focal-plane readout (§2.2). These sections are evident in the Raw Images. All Raw Images, from both E2V and ITL sensors, have the same dimensions: 4608 × 4096 pixels. However, the two types have slightly different numbers of overscan and pre-scan pixels, so the number of exposed pixels also differs. This difference is reflected in the dimensions of the Visit Images (§4.1.2), which have had the prescan and overscan regions removed during data processing.

There are a total of 5,339,820 Raw Images in DP2.

4.1.2. *Visit Images*

Visit Images are the first scientifically useful images to be produced by the processing. They are the result of removing instrument signatures from the raw input images, followed by multiple rounds of background subtraction and photometric and astrometric calibration (see §5). There is one Visit Image for every raw input image that successfully passed through the first stage of the processing. As such, there are up to 180 Visit Images for each of the exposures taken during the observing campaign.

During ISR, the science processing pipeline removes the overscan regions from the input raw images. The Visit Images therefore only contain the pixels that were exposed during the observation: 4072 ×4000 pixels in the case of the ITL detectors, and 4096 ×4004 pixels in the case of the e2v detectors. Unlike raw images, Visit Images consist of three image planes: a science image plane, a variance plane, and a mask plane. The science image is the processed exposed image, whereas the variance image persists the uncertainty associated with the value of the corresponding pixel in the science image. The mask image indicates pixels that are flagged for various different reasons, such as saturated pixels, pixels associated with detected sources, bad pixels, etc. All three image planes contain the same number of pixels in each dimension.

There are a total of 4,988,654 Visit Images in DP2.

4.1.3. *Coadds*

Visit Images covering the same patch of sky (as defined by the skymap) in the same band are combined to form coadd images. Two different types of coadd images are included in DP2: Deep Coadds and Template Coadds, Each type of coadd image are produced via the same general process (§5.3.2): Visit Images are selected by patch, band, and image quality, warped onto a common pixel grid, and coadded. The main differences that distinguish each type of coadd are the selection of the Visit Images, and that Deep Coadds are cell-based in DP2 (see the introduction of this this section), whereas Template Coadds are not. As a result, the structure of the datasets is the same for each type: each coadd contains three planes corresponding to the science image plane (i.e., the coadd), the variance plane, and the mask plane, with all three containing the same number of pixels in each dimension. Up to two different background models are included in the coadd datasets: the background that was subtracted to create the final coadd and, in the case of the *gri* band Deep Coadds, a background that can be added-back to create "Visualization Coadds" which, while not recommended for science, can be useful for visualization. Each coadd also contains a position-dependent PSF model, WCS information, plus various other metadata.

The same skymap was used for each type of coadd in DP2, meaning they all cover the same on-sky area, including some overlap with their neighbouring coadds (see the introduction to this section). A single coadd, irrespective of type, contains 3300 ×3300 equal-sized square pixels, including overlap; each pixel is 0.″2 on each side. There are up to six coadds of each type per patch - one for each of the six LSSTCam bands. While the sky coverage of coadds is similar between types and bands, they are not exactly the same due to the different Visit Image selection functions for each type, and differences in sky coverages in each band.

Table 7. Per-detector PSF-quality selection cuts for inclusion in Deep and Template Coadd images.

Parameter	Description	Numerical cut
PSF FWHM	Same FWHM calculation as the visit-level cut, evaluated per detector rather than averaged over the visit.	< 1.8 arcsec
Median ellipticity residual	Median of $\sqrt{e_1^2 + e_2^2}$ between the PSF model and measured star shapes.[a]	< 0.19
`psfStarScaledDeltaSizeScatter`	Robust scatter of PSF size residuals, scaled by the median PSF size.	< 0.07
`psfTraceRadiusDelta`	Max−min of the PSF trace radius evaluated on a grid across the unmasked detector pixels.	< 4.4 pixels
`psfApFluxDelta`	Max−min of the model PSF aperture flux (normalized) across the unmasked detector.	< 1.6
`psfApCorrSigmaScaledDelta`	Max−min of the PSF aperture correction across the detector, scaled by `psfSigma`.	< 0.15
`starEMedian`	Median star ellipticity, $\sqrt{e_1^2 + e_2^2}$.[a]	< 0.2 (all bands)
`starUnNormalizedEMedian`	Same, computed from unnormalized second moments.[b]	< 2.3 (all bands)
`nPsfStar`	Number of PSF-candidate stars available on the detector; below this the PSF model is not considered well-constrained.	≥ 6 per band

[a] $e_1 = (I_{xx} - I_{yy})/(I_{xx} + I_{yy})$, $e_2 = 2I_{xy}/(I_{xx} + I_{yy})$, where I_{xx}, I_{yy}, I_{xy} are the adaptive second moments of the PSF/star.

[b] Unnormalized equivalent: $\sqrt{(I_{xx} - I_{yy})^2 + (2I_{xy})^2}$, i.e. without dividing by $I_{xx} + I_{yy}$.

The number of Visit Images that contribute to a coadd varies across the patch. This is true even if a patch is fully covered by all the visits that targeted that area because of the presence of chip gaps between neighbouring detectors. The Survey Property maps (see §4.3.1) can be consulted to gain an insight into the total exposure time, and thus the number of contributing Visit Images, at any point within a coadd. Since each coadd image always covers an entire patch, it is not uncommon for regions of a coadd image to have zero Visit Image coverage. This is especially the case on the edges of an observed field. While such regions may contain pixels with apparently valid values (i.e., nonzero or non-`NaNs`), they will be flagged as `NO_DATA` in the mask image plane. It is therefore crucial that the mask image plane is consulted when analysing any type of coadd.

Deep Coadds—Deep Coadds are optimized for depth and reach to the faintest magnitudes of all coadd types. Two basic selection criteria were used to select visits for inclusion in Deep Coadds. The first is seeing-based and selects visits with a mean (over all detectors) PSF FWHM$<1\farcs7$. The second is depth-based and uses the concept of "effective exposure time", which is the exposure time that would be needed under nominal conditions to achieve the same signal-to-noise as the visit in question (Bernstein et al. 2016). The selection uses the transparency component of effective exposure time, and selects visits in which the whole-system transparency was such that the mean (over all detectors) transparency component of the effective exposure time was > 0.75 times the actual exposure time (corresponding to ~ 0.16 mag of extinction relative to nominal conditions). Following the selection of visits, a second round of selection is applied at the detector (i.e., Visit Image) level. These per-detector selection criteria are summarised in Table 7. Together, these selections aim to maximise the depth of the coadd while still ensuring good image quality. Once the selected Visit Images have been warped and coadded, a further coadd-level background is estimated and subtracted from each Deep Coadd during object detection. This background is estimated from the iteratively 3σ-clipped mean of the unmasked pixels (i.e., excluding detected sources) in 128×128 pixel bins, which is then interpolated to full resolution using an Akima spline.

There are a total of 925,460 Deep Coadds in DP2.

Template Coadds—Template Coadds are created, principally, to act as reference images in the process of image subtraction to produce Difference Images (see §4.1.4). When performing image subtraction, the point spread function (PSF) of the reference image is degraded to match that of the science image. The reference image

should therefore have as small a PSF as possible whilst being deep enough to not introduce substantial noise during subtraction (i.e., so that the noise budget is dominated by the single visit image, rather than the template coadd). Template Coadds include only visits with a median (over all detectors) PSF FWHM<1.″4 and a mean (over all detectors) transparency component of the effective exposure time (see *Deep Coadds* above) > 0.75 times the actual exposure time. Of the visits that meet these criteria, the third that cover a given patch in a given band with the smallest median PSF are selected for coaddition. If that third corresponds to fewer than 12 visits (i.e., there are fewer than 36 visits that cover the region in a given band), then the 12 visits with the smallest median PSFs are selected. If there are fewer than 12 visits covering a patch with a given band, then all visits that meet the above PSF FWHM and sky transparency criteria are selected. This strategy optimizes for image quality in regions that are well-covered by visits, while also ensuring good Template Coadd coverage in regions that have poor visit coverage. Following the selection of visits, a second round of selection is applied at the detector level. These per-detector selection criteria are summarised in Table 7.

There are a total of 1,781,170 Template Coadds in DP2.

4.1.4. *Difference Image*

Difference Images are created by subtracting a warped, scaled and PSF-matched Template Coadd from a Visit Image. Because a Visit Image may cover more than one patch or even tract, multiple Template Coadds may also have to be stitched together prior to being subtracted from the Visit Image. The concept behind Difference Imaging is that, following subtraction, only those sources that have either moved or whose flux has changed relative to the template should remain. In practice, however, many spurious sources remain due to unavoidably imprecise template matching. A detection algorithm is used to identify sources in Difference Images and a real/bogus classifier attempts to distinguigh real from spurious sources. The outcome of the real/bogus classifier is included in the DIA Source catalog (§4.2.2).

Like Visit Images, Difference Images also consist of three image planes: a science image plane, a variance plane, and a mask plane (see §4.1.2 for a description of each plane).

There are a total of 3,552,364 Difference Images in DP2, one for each Visit Image.

4.2. *Catalogs*

Catalogs contain data for detections in the images described in §4.1, except the Calibration catalog, which contains reference data obtained from previous surveys. Observatory-produced metadata tables are described in §4.5. Rubin adopts a naming convention in which the term "object" refers to an astrophysical object, whereas the term "source" refers to a single-epoch detection of that object. As such, sources are associated with detections in single-epoch Visit Images and Difference Images, whereas objects are associated with detections in multiepoch coadds. Photometric quantities in catalogs are stored as fluxes rather than magnitudes, with fluxes measured in nanojanskies ($1\ \mathrm{nJy} = 10^{-35}\ \mathrm{W\,m^{-2}\,Hz^{-1}}$). Fluxes are preferred for multiepoch observations, as they can be averaged across epochs, unlike magnitudes. Additionally, flux measurements on difference images (§4.1) are computed relative to a template and can therefore be negative, particularly for faint sources in the presence of noise.

The DP2 catalogs each vary in terms of their specific columns but most contain: one or more unique identification numbers, positional information, multiple types of flux measurements (e.g., aperture fluxes, PSF fluxes, Gaussian fluxes, etc.), and a series of Boolean flags indicating characteristics such as saturation or cosmic-ray contamination for each source or object. The solar system catalogs SSObject and SSSource deviate from this general structure in that they instead contain orbital parameters for all known asteroids.

Where applicable, quantities are prefixed with the band in which they were measured, and all measured properties are reported with their associated 1σ uncertainties. For example, `g_ra` and `g_raErr` in the Object catalog refer to R.A. and its uncertainty, measured in the g band. Fluxes for various apertures are provided together with an uncertainty and a flag, and named in the format `[band_]ap[size]Flux`, where `[size]` is the aperture diameter in pixels, and `[band_]` is used where applicable (e.g., in the object catalog, but not in the source catalog; see §4.2.1 and §4.2.3). For example, `g_ap03Flux`, `g_ap03FluxErr`, `g_ap03Flux_flag` provide the flux, uncertainty and a flag measured within a 3.0 pixel aperture in the g band. Similarly for flux measurements using difference algorithms, e.g. `g_psfFlux` provides the flux derived using the PSF model as a weight function, forced on g band.

All of the catalog types that were released as part of DP1 have been retained in DP2, although in some cases additional columns have been added for DP2; we highlight prominent additional columns in the subsections below. There are also a number of additional catalogs in DP2 that do not have DP1 equivalents. Throughout this section, we provide high-level overviews of the contents of each catalog, rather than listing and de-

scribing each column. Instead, full searchable catalogue schemas, including column descriptions, are provided at https://sdm-schemas.lsst.io/.

All catalogs except the Calibration Reference Catalog are stored in the Qserv database (§7.5.1) and accessible via Table Access Protocol (IVOA standard) (TAP), with an online DP2 catalog schema available at https://sdm-schemas.lsst.io/dp2.html. Catalog data are also accessible via the Data Butler (see §7.2.2).

4.2.1. *Source*

The Source catalog contains the measured properties of all sources that are, prior to deblending, detected at a $> 5\sigma$ significance in each of the Visit Images (§4.1.2). The detections in the Source catalog have all undergone deblending; in the case of blended sources, only the post-deblended sources are included in the Source catalog. It is therefore possible for the Source catalog to contain entries in which none of the different types of photometric measurements (e.g., aperture, PSF, gaussian) have flux-to-error ratios greater than five.

Each entry in the source catalog has a unique identification number (`sourceId`), positional information, various photometric measurements (aperture, PSF, gaussian), shape and extendedness information, and a series of flags. Because each entry in the source catalog is derived from a detection in a single epoch observation taken with a single filter, each entry also has an associated band.

The Source catalog contains data for 18 billion sources in DP2.

4.2.2. *DIA Source*

The DIA Source catalog contains data on sources detected at $> 5\sigma$ significance in the Difference Images, including those associated with known solar system objects. Following the detection and measurement of each DIA source, an attempt was made to associate it with a DIA object (see §4.2.4). If no suitable DIA object was found, then a new entry was created in the DIA Object catalog. So as well as being indexed on a unique `diaSourceId` identification number, each entry in the DIA Source catalog also contains the identification number of its associated DIA object (`diaObjectId`). The detections reported in the DIA Source catalog have not undergone deblending.

The DIA Source catalog contains data for 1 billion DIA sources.

4.2.3. *Object*

The Object catalog contains data on all objects detected with a greater than 5σ significance in the Deep Coadds. With coadd images produced on a per-band basis, a $> 5\sigma$ detection in one or more of the bands will result in an object being included in the Object catalog. For cases where an object is detected at $> 5\sigma$ in more than one band, a crossmatching has been performed between bands to associate an object in one band with its counterpart(s) in the other bands. As such, each entry in the Object catalog associates objects across multiple bands. The objects reported in the Object catalog have also undergone deblending; in the case of blended detections, only the deblended child objects are included in the catalog. As with the Source catalog, the criterion for inclusion in the Object catalog is a $> 5\sigma$ detection in one of the Deep Coadds prior to deblending, yet the positions and fluxes of objects are reported postdeblending. Hence, it is possible for the Object catalog to contain objects whose flux-to-error ratios are less than five in all bands.

In addition to the general information mentioned above (i.e., IDs, positions, fluxes, and flags), the Object catalog also includes basic shape and extendedness information. While they may be included in future Data Releases, no photometric redshifts, Petrosian magnitudes (Petrosian 1976), proper motions or periodicity information are included in the DP2 object catalogs.

The Object catalog contains data for 790 million objects in DP2.

4.2.4. *DIA Object*

The DIA Object catalog contains the astrophysical objects that the DIA sources are associated with (i.e., the DIA objects). Only non-solar-system objects are contained in the DIA Object catalog. Solar system objects are recorded in the Solar System Object catalog (§4.2.10). Once a DIA source has been detected and measured, the DIA Object and Solar System Object catalogs are searched for objects to associate it with. If no association is found, a new entry in the DIA object catalog is created and the DIA source is associated to it. An attempt has also been made to associate DIA objects across multiple bands, meaning each entry in the DIA Object catalog may contain measurements from multiple bands. Because they ultimately stem from detections in Difference Images, DIA objects are typically transient or variable. Due to this variable nature of DIA objects, the DIA Object catalog contains summary statistics of their fluxes, such as the mean and standard deviation over multiple epochs. For single epoch measurements of the properties of DIA objects, users should refer to either the Forced Source at DIA Object catalog (§4.2.6) or the DIA Source catalog for single-epoch flux measurements of DIA objects, both of which contain the DIA object identification number (`diaObjectId`) for easy look-up.

The DIA Object catalog contains data for 230 million DIA objects.

4.2.5. *Forced Source*

The Forced Source catalog contains forced PSF photometry measurements performed on both Difference Images (i.e., the `psfDiffFlux` column) and Visit Images (i.e., the `psfFlux` column). Forced photometry is performed on each Visit Image and Difference Image at the positions of all the objects in the Object catalog that is covered by that image. The information in the Forced Source therefore enables the assessment of the time variability of the photometry of each object in the Object catalog. Unless there is strong justification to do otherwise, we recommend using the forced photometry measured on the Difference Images (i.e., `psfDiffFlux`) rather than the Visit Images (i.e., `psfFlux`). Changes in the PSF between visits can make flux from even non-varying neighboring sources appear as variability in `psfFlux`, whereas such sources are subtracted out in the Difference Images.

In addition to forced photometry PSF fluxes, a number of Boolean flags are also included in the Forced Source catalog.

The Forced Source catalog contains a total of 45 billion entries across the 790 million unique objects.

4.2.6. *Forced Source at DIA Object*

The Forced Source at DIA Object catalog is equivalent to the Forced Source catalog, but contains forced photometry measurements performed on each Visit Image and Difference Image at the positions of all the DIA objects in the DIA Object catalog.

The Forced Source at DIA Object catalog contains a total of 24 billion entries across the 230 million unique DIA objects.

4.2.7. *Isolated Star Stellar Motions*

The Isolated Star Stellar Motions catalog contains the proper motions and parallaxes of isolated stars. Rubin uses a specific definition for an isolated star. Candidate isolated stars must be point-like (i.e., unresolved) detections with a signal-to-noise ratio of between 8 and 1000, and must not be flagged as blended. All such candidates within a given tract are spatially matched using a 1″matching radius across all visits and bands covering said tract. Should the average position of a candidate have no other candidate within 2″, then it is considered to be an isolated star.

The Isolated Star Stellar Motions catalog provides a best-fit position (`ra, dec`), proper motion in mas yr^{-1} (`raPM, decPM`), and parallax in mas at a reference epoch; the `raPM` parameter already includes the cos(decl.) factor. The reference epoch in Modified Julian Date (MJD) is included as a column (`epoch`) in the catalog. An isolated star's proper motion and parallax are calculated by fitting its position across all per-visit detections, plus any detection in our calibration reference catalog (§4.2.12) that is matched to within 0.″1 of the star's average position. The catalog also contains 1σ uncertainties for each each of the five aforementioned positional parameters (five columns: `raErr, decErr, raPMErr, decPMErr, parallaxErr`), plus a further ten columns containing off-diagonal covariance parameters (`*_*_Cov`, e.g, `ra_raPM_Cov`, `raPM_parallax_Cov`).

All isolated stars have an entry in the Isolated Star Stellar Motion catalog, although no attempt is made to fit the time-varying positions of stars that are detected in fewer than 3 epochs and which do not have a matched counterpart in the reference catalog. In such cases the positional information, including uncertainties, are reported as `NaN`. Every entry in the Isolated Star Stellar Motion catalog has a unique identification number and, in cases where a match to the reference catalog has been identified, the unique identification of the matched reference source.

The Isolated Star Stellar Motions catalog contains 320 million entries.

4.2.8. *Shear Object*

The Shear Object catalog contains object shape measurements in five different shear steps or "counterfactual" images: one in which no shear (`ns`) has been applied, and four in which an artificial linear shear has been introduced to the image (`1p, 1m, 2p, 2m`). The artificial shears are in four different directions: 1p is along the image x-direction, 1m is along the y-direction, 2p is along the 45 degree diagonal, 2m is along the 135 degree diagonal.

Two measurement algorithms (*gauss* and *pgauss*) are run on each of the five different counterfactuals, one for measuring shapes and the other for fluxes. The first approach is to fit a single Two-dimensional (2D) gaussian jointly across the r-, i-, and z-band deep coadd images, with the ellipticity components in the image x-y (`gauss_g1`) and diagonal (`gauss_g2`) directions reported in the catalog. The second approach measures 2D gaussian-weighted moments in the g, r, i, and z bands (e.g., `r_pgaussFlux`). Like the Gaussian Aperture and PSF (GAaP) fluxes in the Object catalog, these are designed to provide consistent colors across bands rather than total fluxes; the second approach does not provide ellipticity measurements. There is one row per object detected in each counterfactual, meaning a given object could result in up to five rows, depending on how

many of the counterfactuals it is detected in at a significance of $> 5\sigma$.[34]

In addition to the gaussian shape measurements and their covariances, the Shear Object table also contains object positional information, fluxes measured from both the gaussian fits and the gaussian moments (g, r, i and z bands), PSF shapes reported in the same x-y and diagonal directions as objects, a range of different flags. Each entry in the Shear Object catalog also has a unique identification number (`shearObjectId`), that uniquely identifies a given detection in a given counterfactual. The Shear Object catalog contains 2.5 billion rows.

4.2.9. *Solar System Source*

The Solar System Sources catalog is derived from the Rubin observations. It contains one row per associated detection and links measured DIA Source photometry and astrometry to a predicted object in the Solar System Object catalog. As entries in the Solar System Source catalog stem from the DIA Source catalog, they have all been detected at $> 5\sigma$ significance in at least one band.

The `diaSourceId` can be used to identify corresponding detections in the DIA Source catalog; each `diaSourceId` occurs at most once in the Solar System Source catalog. The `ssObjectId` or `designation` can be used to identify corresponding entries in the object-level catalogs (i.e., Solar System Object, Minor Planet Center (MPC) Orbit, Current Identifications, and Numbered Identifications).

Each entry in the Solar System Source catalog carries the measured position and the Sorcha (Merritt et al. 2025; Holman et al. 2025) ephemeris position, the measured PSF photometry and band, the predicted apparent V-band magnitude, phase angle and solar elongation, heliocentric and topocentric distances, sky-plane and radial rates, and Cartesian position and velocity components. R.A. and decl. are in degrees in the ICRS frame; the observation epoch `DIA_midpointMjdTai` is an MJD on the International Atomic Time (TAI) time scale. State vectors are equatorial and evaluated at the light-emission (retarded) time. The `ephVmag` field is an approximate predicted Johnson V magnitude derived from the input catalog photometric parameters. It is not a prediction in the Rubin band of the associated detection, so comparisons to the measured *ugrizy* photometry include both object color and phase-model differences.

The Solar System Source catalog contains 8.1 million rows for 299,343 distinct objects.

4.2.10. *Solar System Object*

The Solar System Object catalog is also derived from Rubin observations. It contains one row per object and summarizes its observations, per-band photometric properties, including absolute magnitudes, and selected orbit-derived quantities. The Solar System Object catalog contains 299,343 entries; one entry for each distinct object. There are 17,528 objects with a single associated detection, 31,873 with at most two, 212,181 with at least ten, and 7,293 with at least 100. The main belt dominates the distance and rate distributions, but the catalog includes 35 objects observed beyond 50 astronomical unit (au), 765 source records at topocentric distances below 0.1 au, and 2,105 records with predicted rates above 1 deg day^{-1}. The ecliptic-latitude range of the objects, $-69.4°$ to $+13.3°$, primarily reflects the southern DP2 footprint rather than the intrinsic solar system distribution.

4.2.11. *MPC Orbit, Current Identifications, Numbered Identifications*

The MPC Orbit, Current Identifications, and Numbered Identifications catalogs are each auxiliary snapshots of tables in the MPC PostgreSQL database. They provide reference catalogs for orbital information and object cross-identification, retaining the upstream MPC semantics rather than being recomputed from DP2 observations. Their field definitions are documented in the MPC schema for the PostgreSQL tables replicated through the Small Bodies Node.[35] The MPC Orbits catalog records the MPC orbit solution and associated fit information for designated objects. The Current Identifications catalog identifies primary and secondary provisional designations that refer to the same object, and the Numbered Identifications catalog connects an object's permanent number to its primary provisional designation. The former includes both self-mappings ($A = A$) and links among aliases ($A = B$); the latter also carries International Astronomical Union (IAU) naming information when available. The MPC Orbits catalog available at ingestion did not contain comets. Because the predictions used for association were generated from that snapshot, no comets were available to be associated with DP2 detections and none occur in the delivered Solar System Source or Solar System Object catalogs.

[34] It is feasible for a marginal 5σ object in one counterfactual to be just below the 5σ detection threshold in another counterfactual.

[35] https://docs.minorplanetcenter.net/mpc-ops-docs/data-and-services/replicated-tables-schema/

In the Solar System Source, Solar System Object, and MPC Orbits catalogs, `designation` denotes the unpacked primary provisional designation of the object. This was added to the delivered NSF Mathematical and Physical Sciences directorate (MPS) Orbits copy as an end-user convenience; there it is identical, row by row, to `unpacked_primary_provisional_designation`. It therefore provides a readable and direct join key among the Solar System Source, Solar System Object, MPC Orbits catalogs, without requiring users to unpack MPC designations. For query performance, users should nevertheless prever to use integer identifiers such as `ssObjectId` to link catalogs wherever those identifiers are available, and use `designation` for readability or when an appropriate integer identifier is unavailable.

4.2.12. *Calibration Reference Catalog*

The Calibration Reference Catalog is used to perform astrometric and photometric calibration. It is a whole-sky catalog built specifically for LSST by combining data from multiple prior reference catalogs, since no single existing catalog had both the depth and coverage needed to calibrate LSST data. It contains only stellar sources. Full details on how the Calibration Reference Catalog was built are provided in (Ferguson et al. 2025).[36] The same reference catalog was used for calibrating DP2 as for DP1, but we prefer to include the same summary here over only providing a reference to that previous data preview.

For the *grizy* bands, the input catalogs were (in order of decreasing priority): Dark Energy Survey (DES) Year 6 Calibration Stars (Rykoff et al. 2023); Gaia B or R Photometry (Gaia) (XP) synthetic magnitudes (Gaia Collaboration et al. 2023); the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS) 3PI survey (Chambers et al. 2016); Data Release 2 of the Skymapper survey (Onken et al. 2019); and Data Release 4 of the VLT Survey Telescope (VST) Asteroid Terrestrial-impact Last Alert System (ATLAS) survey (Shanks et al. 2015). For the u band, the input catalogs were (in order of decreasing priority): standard stars from Sloan Digital Sky Survey (SDSS) Data Release 16 (Ahumada et al. 2020); Gaia XP Synthetic Magnitudes (Gaia Collaboration et al. 2023); and synthetic magnitudes generated using Stellar Locus Regression (SLR), which estimates the u-band flux from the g-band flux and g - r colors. These SLR estimates were used to boost the number of u-band reference sources, as otherwise the source density from the u-band input catalogs is too low to be useful for the LSST.

Throughout, the Calibration Reference Catalog uses the DES bandpasses for the *grizy* bands and the SDSS bandpass for the u-band; color transformations derived from high-quality sources were used to convert fluxes from the various input catalogs (some of which did not use the DES or SDSS bandpasses) to the respective bandpasses. All sources from the input catalogs are matched to Gaia Data Release 3 (DR3) sources for robust astrometric information, selecting only isolated sources (i.e., no neighbors within $1''$).

After collating the input catalogs and transforming the fluxes to the standard DES and SDSS bandpasses, the catalog was used to identify sources within a specific region of the sky. This process generated a set of standard columns containing positional and flux information, along with their associated uncertainties.

4.3. *Maps*

4.3.1. *Survey Property Maps*

Survey Property maps capture a number of position-dependent quantities such as observing conditions and exposure time. For each band, a given map encodes the spatial distribution of one such quantity, built up from the images that are combined to create the Deep Coadds. The construction of the Survey Property maps is initially performed at the tract-level, then per-tract maps are combined into a single consolidated map; only these consolidated maps are provided. Maps displaying different summary statistics, such as minimum, maximum, mean, standard deviation, of the various properties are provided. Each sky position corresponds to a spatial pixel in the Hierarchical Equal-Area iso-Latitude Pixelisation (HEALPix)(Górski et al. 2005) grid.

There are fourteen Survey Property maps in DP2. Among the quantities mapped are the total exposure time; three observation-epoch maps (earliest, mean, and latest epoch); the PSF size and shape, the latter given as separate e^1 and e^2 components; the PSF magnitude limit; sky background and noise; astrometric offsets in R.A. and decl.; and, for the wavelength-dependent atmospheric Differential Chromatic Refraction (DCR) effect, PSF distortion again split into e^1 and e^2 components.

The names of the maps follow the pattern `deepCoadd_<PROPERTY>_consolidated_map_<STATISTIC>`—so `deepCoadd_exposure_time_consolidated_map_sum`, for instance, gives the summed exposure time, in seconds, at each sky position. The maps

[36] In (Ferguson et al. 2025), the calibration reference catalog is referred to as "The Monster".

themselves are stored in `HealSparse`[37] format, and Survey Property maps are only accessible through the Data Butler (§7.2.2).

4.3.2. *HiPS Maps*

Hierarchical Progressive Survey (IVOA standard) (HiPS) maps, offer an interactive way to explore seamless, multiband tiles of the sky regions covered by DP2, allowing for smooth panning and zooming. DP2 provides multiband HiPS images created by combining data from individual bands of Deep Coadd images, using an improved version (N. B. Lust et al. in preparation) of the algorithm presented in (Lupton et al. 2004). These images are false-color representations generated using various filter combinations for the red, green, and blue channels.

HiPS maps are only accessible through the HiPS viewer in the RSP Portal (§7.3) and cannot be accessed via the Data Butler (§7.2.2).

4.4. *Ancillary Data Products*

In addition to the science images and catalog data products described above, there are a number of ancillary data products included in DP2. While not science-level data products *per se*, their inclusion in DP2 will help to facilitate science.

4.4.1. *Skymap*

The DP2 skymap, `lsst_cells_v2`, is also included as a data product in DP2. The skymap data product can be used, for example, to identify the tract and patch that contains a given sky coordinate, or obtain the coordinates of the corners of a tract or patch.

A description of the `lsst_cells_v2`, in terms of the numbers and sizes of its tracts and patches, is given at the start of this section.

4.4.2. *Coadd Patches*

The Coadd Patches table lists, for each patch in the skymap covered by DP2, its tract and patch IDs, central coordinates, and the region of sky it covers. It can be used to identify which patches, and hence which coadds, cover a given position on the sky. It is only available via the Qserv database (§7.5.1).

4.4.3. *Standard Passbands*

Figure 2 shows the full-system throughput of the six LSSTCam filters. The corresponding transmission curves are provided as a DP2 data product. These datasets tabulate the full-system transmission of the six LSSTCam filters as a function of wavelength and were used as a reference for the LSSTCam DP2 photometry. The Standard Passband dataset is keyed by band and is stored in Astropy table format.

4.4.4. *Object Scarlet Models*

The Object Scarlet Models dataset stores the scarlet models fit to each deblended detection in the Deep Coadd images. Each model captures an object's spatial profile (morphology) together with its per-band flux values, as assigned by the deblender when separating it from its blended neighbours. When combined with the corresponding Object catalog and Deep Coadd image, this dataset enables the reconstruction of the deblend model to see how a specific detection was deblended. There is one Object Scarlet Models dataset per patch that has coverage in DP2, and they are only available via the Data Butler.

4.4.5. *Deep Coadd Input Summary*

The Deep Coadd Input Summary table contains the visit and detector IDs of all the Visit Images that were selected and coadded to produce each Deep Coadd image. Also included in the table are the relative weight of each Visit Image, inversely proportional to its variance, and the number of pixels it contributed to a Deep Coadd, since a given Visit Image may only partially cover a patch. There is a single Deep Coadd Input Summary table for the whole of DP2, containing one row per contributing Visit Image per patch (i.e., one row per unique combination of tract, patch, visit, and detector). It is only available via the Data Butler.

4.4.6. *Compare Warp Artifact Mask*

The Compare Warp Artifact Masks are the mask planes of the individual warped Visit Images, in which pixels affected by transient artifacts, such as satellite trails and cosmic rays, are flagged in the `CLIPPED` plane and excluded from the Deep Coadds. They are generated by the artifact rejection procedure described in §5.3.2. There is one Compare Warp Artifact Mask per Visit per patch, and they are only available via the Data Butler.

4.4.7. *Camera*

The Camera describes the geometry of LSSTCam, i.e., the positions and orientations of the detectors in the focal plane, the layout and readout geometry of each detector's amplifiers, and a nominal model of the optical distortion. It is used throughout processing to map pixels on each detector to positions in the focal plane, for

[37] A sparse HEALPix representation that efficiently encodes data values on the celestial sphere. https://healsparse.readthedocs.io

example when constructing the initial World Coordinate System (WCS) of each exposure. While of limited use as a standalone data product, the Camera is a required input to a number of processing steps. It is therefore largely provided to enable users to rerun selected processing tasks on subsets of DP2. There is a single Camera for DP2, and it is only available via the Data Butler.

4.4.8. *FGCM Lookup Table*

The FGCM Lookup Table contains pre-computed integrals of the combined atmospheric and instrumental throughput in each band over the expected range of atmospheric conditions. It is generated from the MODTRAN atmosphere model and the measured filter and detector throughputs (§5.2.3), in the same processing step that produces the Standard Passbands (§4.4.3), and it allows the Forward Global Calibration Method (FGCM) photometric calibration to rapidly compute throughput and chromatic corrections for stars of different colors. There is a single FGCM Lookup Table for DP2, and it is only available via the Data Butler.

4.4.9. *Pretrained Model Package*

The Pretrained Model Package contains the architecture and trained weights of the convolutional neural network used to assign a reliability (i.e., real/bogus) score to each DIA source (§5.4.2; Acero-Cuellar et al. 2026). The same model was used for both DP1 and DP2. There is a single Pretrained Model Package for DP2, and it is only available via the Data Butler.

4.4.10. *Run Provenance*

The Run Provenance dataset information regarding the DP2 processing itself: the inputs, outputs, configurations and logs of every processing task. It is only available via the Data Butler.

4.4.11. *Calibration datasets*

A number of datasets used during ISR and for calibration are also included in DP2. We anticipate that the vast majority of users will not need to use these datasets, but we include them for transparency and completeness. Below, we provide a brief description of each; more detailed descriptions of how they are generated and used can be found in Plazas Malagón et al. (2025).

Bias—The final bias frames, one per detector, that are generated by taking the sigma-clipped per-pixel mean of many (typically around 20) zero-second exposures of the detectors. They are used during ISR to remove the fixed-pattern offset structure inherent to the readout electronics.

Dark—The final dark frames, one per detector, that are generated by taking the sigma-clipped per-pixel mean of many nonzero-length dark images taken with the shutter closed and spanning a range of different integration times. They are used during ISR to remove the contributions of dark current and light leaks.

Flat—The final flat frames, one per band per detector, that are generated by combining many exposures of the Rubin in-dome flat field system (see §5.1.1). They are used during ISR to correct for pixel-to-pixel differences in detector sensitivity.

Brighter Fatter Kernel—The Brighter Fatter Effect is caused by the build-up of charge deflecting subsequent photo-electrons into neighbouring pixels, resulting in brighter objects appearing larger than they otherwise would. The Brighter Fatter Kernel (see §5.1.1) is used to correct this effect during ISR.

Crosstalk—The cross talk matrices, one per detector, that quantify how much of the signal read out through each amplifier leaks into the other amplifiers of the same detector. They are used during ISR to remove the faint artifacts that bright sources in one amplifier segment induce in the others.

Charge Transfer Inefficiency—The deferred-charge calibrations, one per detector, that contain a per-amplifier model of serial charge transfer inefficiency (CTI): the incomplete transfer of charge towards the readout amplifiers, which causes some charge to be trapped and later released as faint trails behind sources. They are used during ISR to remove these trails.

Defects—The defect mask, one per detector, that flags pixels and columns as unusable for science. Defects are identified as outliers in the pixel value distributions of dark and flat-field images, using either individual or combined frames. The individual dark images used in this process are taken at a range of different integration times, in order to also identify dynamic defects (e.g., bad columns) that only manifest at certain exposure durations.

Linearizer—The linearity corrections, one per detector, that contain a per-amplifier model of the nonlinear response of the signal chain, derived from pairs of flat-field images taken over a range of flux levels. They are used during ISR to correct each pixel value for this nonlinearity.

Photon Transfer Curve—The photon transfer curves (PTCs), one per detector, that describe how the variance and the covariances between neighboring pixels in each amplifier vary with signal level, measured from

pairs of flat-field images taken over a range of flux levels. They are used to measure the gain of each amplifier, which is used during ISR to convert pixel values from instrumental units to electrons, and the pixel covariances from which the Brighter Fatter Kernel is derived.

4.5. *Metadata*

4.5.1. *Visit catalog*

The Visit catalog contains technical data for each visit collected at the time of observation, such as telescope pointing, camera rotation, airmass, exposure start and end time, and total exposure time. Some of the information contained within the Visit catalog is also contained in the Visit Detector catalog described below (e.g., exposure time), although the latter also includes information produced by the processing pipelines at the per-detector level, such as the PSF size and limiting magnitudes of a given Visit Image.

The Visit catalog contains 28,698 entries.

4.5.2. *Visit Detector catalog*

The Visit Detector catalog contains summary information for all Visit Images in DP2. In addition to technical information, such as the on-sky coordinates of the central pixel and measured pixel scale, the Visit Detector catalog contains a range of data quality measurements, such as whole-image summary statistics for the PSF size, zero-point and 5σ depth, sky background, sky noise, and the quality of astrometric solution. It provides an efficient method to access the properties of a Visit Image without needing to access the image data. When combined with the data contained in the Visit table described above, it provides a full picture of the telescope pointing and sky conditions at the time of observation.

The Visit Detector catalog contains 5.2 million entries.

4.5.3. *Visit Summary*

The Visit Summary table contains the same information as the Visit Detector catalog, but only for a single visit. There is one Visit Summary table per visit (meaning the Visit Detector table is the collation of all the Visit Summary tables), and they are only available via the Data Butler.

5. DATA RELEASE PROCESSING

Data Release Processing (DRP) is the data pipeline that produces the calibrated images, detection catalogs, and derived data products defined in section §4, using version v30.0.8 of the LSST Science Pipelines (Rubin Observatory Science Pipelines Developers et al. 2026).

Data Release Processing (DRP) is the processing mode that produces the calibrated images, catalogs, and derived data products described in §4. For DP2, DRP used version v30.0.8 of the LSST Science Pipelines (Rubin Observatory Science Pipelines Developers et al. 2026). DP2 was processed at the USDF at the SLAC National Accelerator Laboratory, with pilot runs executed at the National Institute of Nuclear and Particle Physics (IN2P3) computing center as part of the ramp-up toward Rubin's multi-site distributed data processing.

The processing campaign consists of four major stages: (1) single-frame processing, (2) calibration, (3) multi-visit processing, (4) variable and transient source processing. Data quality and key metrics for verification and validation are actively monitored during each stage, and comprehensive verification and validation assessments are completed before processing the next stage.

This section describes the high-level algorithms, configurations, and verification metrics for each stage, including changes and improvements from DP1 (Vera C. Rubin Observatory Team et al. 2026).

5.1. *Single-Visit Processing*

5.1.1. *Calibration & Instrument Signature Removal (ISR)*

The first task in data processing begins with a raw image. This image is in discretized analog-to-digital units (ADU, or counts) and contains detector artifacts imprinted by the physics of the detector and readout hardware. This step aims to remove any artifacts that are not of astrophysical origin and produce science-ready, single-visit exposures in floating-point electron units with complete pixel masks.

The ISR pipeline is built on a physical model of the detectors and readout electronics, in which each artifact is tied to a specific point in the photon-transfer chain and is corrected using calibrations from dedicated campaigns conducted during the day or on cloudy nights. Figure 6 shows this model of the detector and readout components and the effects each imparts on the measured image. ISR works backward through the chain, undoing these effects in the reverse of the order in which they were imprinted. Detector response and performance vary across the focal plane (see §2.2), and each detector must be calibrated individually.

A detailed description of the DP2 ISR procedures and the calibration products they consume can be found in (Broughton et al. 2026).

The ISR processing pipeline for DP2 performs the following steps, in order: (1) conversion of the raw integers to floating point, (2) Analogue-to-Digital Unit (ADU) dithering to mitigate quantization effects, (3) serial overscan subtraction, (4) saturated- and suspect-pixel masking, (5) gain (electrons/ADU) correction, (6)

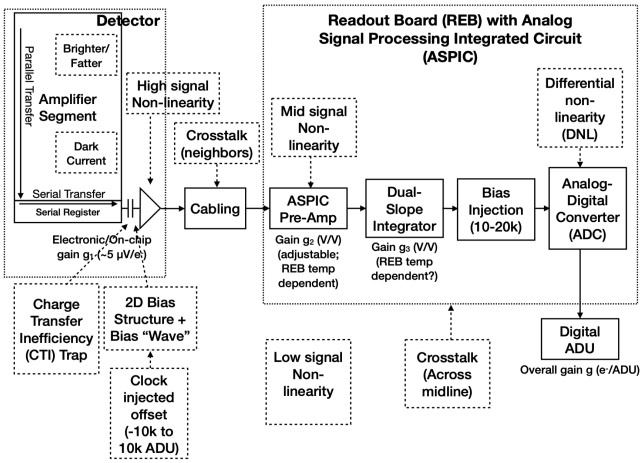


Figure 6. The model of the detector and Readout Electronics Board (REB) components, labeled with the effects that they impart on signal.

cross talk correction, (7) parallel overscan subtraction, (8) nonlinearity correction, (9) deferred-charge correction, (10) vendor-specific masking (E2V edge bleed; the ITL "dip," saturation sag, and edge bleed), (11) bias subtraction, (12) dark subtraction, (13) defect masking, (14) saturated pixel footprint widening, (15) brighter-fatter correction, (16) variance plane estimation, (17) flat fielding, (18) interpolation over masked pixels, (19) amplifier-offset corrections.

Compared to DP1 (Vera C. Rubin Observatory Team 2026), the ISR pipeline has been substantially upgraded, including improved sensor diagnostics and characterization, an entirely new brighter-fatter correction based on Astier & Regnault (2023) and Broughton et al. (2024), and an all-new flat-field production pipeline that corrects for gradients and mis-centering of the flat-field illumination using new calibration hardware. We generate flats using a new set of 10 LEDs spanning the full sensitivity range of the LSSTCam sensors (one in u, two each in *griz*, and a new 1050 nm Light-Emitting Diode (LED) shared between y and z). For the bands covered by two LEDs, we combine separate flats for each LED (for the blue and red side of the band) into a so-called "anaglyph" flat, with a weighting that approximates a flat spectral energy distribution across the band.

Together, these upgrades allow ISR to deliver uniform, science-ready exposures across all LSSTCam science sensors for DP2.

5.1.2. *Background Subtraction*

Background subtraction in `CalibrateImageTask` estimates and removes large-scale background signals from science imaging prior to source detection and measurement. Such signals include sky brightness from airglow, moonlight, scattered light, zodiacal light, and diffuse astrophysical emission. The background is estimated twice: an initial model (`psf_subtract_background`) is subtracted before PSF characterization, and a second model (`star_background`) is fit on the updated source-masked image before the final detection pass. Each post-ISR image is divided into 128×128 pixel superpixels, and the iterative 3σ-clipped mean of unmasked pixels in each bin is fit with a sixth-order two-dimensional Chebyshev polynomial, evaluated at native pixel resolution via Akima spline interpolation. Before fitting `star_background`, the `DETECTED` mask plane is dilated by up to 10 pixels to suppress source-wing flux from leaking into background bins. After `star_background` subtraction, a zeroth-order pedestal correction is estimated iteratively over bin sizes starting at 32×32 pixels, doubling each step until the cumulative pedestal level changes by less than 5% or 0.5 counts between steps. The `star_background` model and pedestal corrections together constitute the full background model.

Compared to DP1 (NSF-DOE Vera C. Rubin Observatory 2025), several changes were made to background subtraction for DP2. The inline re-estimation performed after each detection pass in DP1 (`reEstimateBackground = True`) has been replaced by the dedicated `star_background` subtask. In addition to source mask dilation and a pedestal correction, other new features for DP2 include adaptive threshold determination and diffraction spike masking. The per-step illumination correction used in DP1 has been disabled for DP2. Following background subtraction, the median and standard deviation of all unmasked pixels and of *sky source* fluxes are recorded in the task metadata as diagnostics of background subtraction quality.

5.2. *Calibration*

Calibration of the data consists of estimating and modeling of the point-spread function (PSF), photometric calibration, and astrometric calibration. Calibration is performed using repeatedly detected isolated sources in visit images, detected using a 5σ threshold on the post-ISR images with the background subtracted. Initial astrometric and photometric solutions are derived using only the Calibration Reference Catalog (§4.2.12). Approximate solutions for source positions, fluxes, and PSF shapes are iteratively refined into the final astrometric, photometric, and PSF models described in the following sections.

5.2.1. *PSF Modeling*

PSF modeling in DP2 uses two distinct approaches, as described in the (Rubin Observatory Science Pipelines

Developers et al. 2026) and it is similar to what was originally described in DP1. The first model, which is faster but less precise, is fit with PSFEx (Bertin 2011), while the final model, which is slower but more accurate, is fit with Piff (Jarvis et al. 2021). Broadly speaking, PSFEx and Piff both represent the PSF on a pixel basis and interpolate its parameters across a single CCD using two-dimensional polynomials. However, their technical implementations differ (see Bertin et al. and Jarvis et al. 2021 for more details). Previous comparisons have shown that Piff outperforms PSFEx by delivering an unbiased reconstruction of the PSF size (see Zuntz et al. 2018 and Jarvis et al. 2021), working in sky coordinates, including a color dependency, and supporting alternative models and interpolation schemes (see Schutt et al. 2025 for more details).

In the initial PSF modeling pass, PSFEx is run with a pixel basis and a second-order polynomial interpolation on each CCD. For the final modeling, Piff uses its `PixelGrid` with a fourth-order polynomial interpolation per CCD, except in the u-band, where star counts are too low to constrain a fourth-order fit, so a second-order polynomial is used instead. Both fits are performed in pixel coordinates.

5.2.2. *Astrometric Calibration*

The final astrometric calibration of the data follows the same method as for DP1, but with some key additions. As described in (NSF-DOE Vera C. Rubin Observatory 2025), a joint calibration is performed on all visits in a given band overlapping a given area. By associating all the isolated point sources shared among the overlapping visits, we can refine the astrometric solution beyond the level achieved in single-frame processing using only a reference catalog (Saunders 2024). The full details of the method are described in Saunders (2024).

Unlike in DP1, the area we use for individual fits of the astrometric solution is an order-3 healpix, instead of a tract. This is done to reduce repetitive calculations, since the small size of tracts compared to the LSSTCam field of view means that neighboring tracts often share identical overlapping visits, which would lead to doing many identical fits of the astrometric calibration.

As in DP1, the astrometric model consists of a static map from pixel space to an intermediate frame, which is followed by a per-visit map from the intermediate frame to the plane tangent to the telescope boresight, and finally a deterministic mapping from the tangent plane to the sky. This is done by the `gbdes` package (Bernstein et al. 2017), as described in Saunders (2024). The proper motion and parallax of the objects used in the calibration are fit as part of the solution, but we do not yet include differential chromatic refraction or lateral color. These will be included in future Data Releases during Rubin Operations.

The main difference from the DP1 astrometric calibration is an additional step that models the effect of atmospheric turbulence on the astrometry. The per-visit part of the astrometric model already described captures large-scale variation between visits, which includes some atmospheric effects. However, higher order oscillations are visible in the astrometric residuals, and these can be attributed to atmospheric turbulence (Saunders 2025). As described in Fortino et al. (2021) and Léget et al. (2021), the wave-like astrometric effects caused by turbulence can be modeled by a Gaussian process. In DP2, we do this by taking the astrometric residuals remaining after the `gbdes` fit, grouping them by visit, then modeling them with the Gaussian Process package `treegp`. We then use this Gaussian Processes model to predict the effects of turbulence on a grid. Finally, this grid is turned into a two-dimensional spline that can use the AST component `SplineMap` so that it can be concatenated with the `gbdes` model. By including this model for the higher-order atmospheric turbulence, we achieve a significant improvement in the astrometric calculation.

A corollary product of the `gbdes` astrometric calibration is the catalog of proper motions and parallaxes of the sources used in the calibration. Since the `gbdes` fit is done per-band, and then improved by the turbulence fit, we perform an additional final fit of the proper motion and parallax, using the final astrometric model and all available bands. The resulting catalogs are included in the DP2 Data Release. These catalogs only include isolated point sources, but future Data Releases will be extended to include proper motion and parallax for all objects.

5.2.3. *Photometric Calibration*

Photometric calibration of the DP2 dataset is based on the FGCM (Burke et al. 2018) adapted for the LSST Science Pipelines (Aihara et al. 2022; Fagrelius & Rykoff 2025). We used the FGCM to calibrate the full DP2 dataset with a forward model that uses a parameterized model of the atmosphere as a function of airmass along with a model of the instrument throughput as a function of wavelength. The FGCM process typically begins with measurements of the instrumental throughput, including the mirrors, filters, and detectors. For the LSSTCam filter throughputs, we start with the filter scans provided by the manufacturer, Matereon. These were taken in a grid with 159 positions, at 5 incidence angles from 13.8 to 23 degrees, at a wavelength resolution of 1 nm. We performed an averaging procedure to get the

approximate appropriate throughput for the LSSTCam beam as it intersects the filter at the central location of each detector. For the LSSTCam detector quantum efficiency (QE) curves, we start with per-raft measurements of the QE performed at Brookhaven National Laboratory (BNL) and SLAC at 10 to 25 wavelength steps. These preliminary scans were then normalized relative to each other at 750 nm. However, some per-raft systematic differences are still seen, particularly in the u- and y-bands. In the future we will use the in-dome tunable laser system and the CBP for improved filter and detector throughputs.

The FGCM atmosphere model, provided by MODTRAN (Berk et al. 1999), generates a look-up table of atmospheric throughput as a function of zenith distance at Cerro Pachón, accounting for absorption and scattering by molecular constituents of the atmosphere (including O2 and O3), absorption by water vapor, and Mie scattering by airborne aerosol particulates. Night-to-night atmospheric variation is constrained by minimizing the scatter among repeated "compensated aperture flux" measurements of stars with Signal to Noise Ratio (SNR) greater than 10, where these compensated fluxes include a local background subtraction to remove the effect of background offsets. All six bands (*ugrizy*) are fit simultaneously, with no gray (achromatic) terms allowed beyond a linear term for mirror reflectance degradation and dust accumulation on the lenses; over the 250 days spanned by the DP2 dataset, we find no evidence for a significant loss of throughput. To anchor the overall system throughput and set the "absolute" calibration, the fit is additionally constrained using a subset of stars drawn from the Calibration Reference Catalog[38].

5.3. *Multi-Visit Processing*

5.3.1. *Input Data Selection*

As described in §2.2, detectors 0, 20, 27, 65, 123, 161, 168, 188 are excluded from the Deep Coadd images as they are strongly affected by vignetting. Detector 122 was turned off for the entire DP2 data-taking campaign and is therefore absent from all single-visit and coadded data products. Only exposures with mean seeing better than 1.″7 FWHM and transparency better than 0.75 are included in the Deep Coadd images. Exposures with poor PSF modeling, identified using internal diagnostics and algorithm processing flags, are also excluded from the Deep Coadd images to prevent contamination of the coadds with unreliable PSF estimates.

[38] Referred to as "The Monster" (§4.2.12; see Ferguson et al. 2025).

5.3.2. *Coaddition*

The single Visit Images that pass the selection criteria above are warped onto the coadd pixel grid using Lanczos3 kernel and scaled to have units of `nJy` (zero-point of 31.4). Artifacts are identified and masked in a manner similar to DP1 before coaddition. We apply the artifact rejection procedure described in (AlSayyad 2018), which identifies and masks features such as satellite trails, optical ghosts, and cosmic rays. It operates on a time series of PSF-matched images resampled onto a coadd pixel grid ("warps") and leverages their temporal behavior to distinguish persistent astrophysical sources from transient artifacts. Artifact rejection uses both direct (where no PSF matching is performed) and PSF-matched warps, homogenized to a standard PSF of 1.″8 FWHM, broadly consistent with the 1.″7 FWHM seeing threshold used in data screening. A sigma-clipped mean of the PSF-matched warps serves as a static sky model, against which individual warps are differenced to identify significant positive and negative residuals. Candidate artifact regions are classified as transient if they appear in fewer than a small percentage of the total number of exposures, with the threshold based on the number of visits. Adequate artifact rejection requires a minimum of three overlapping images.

In DP2, the Deep Coadd images are cell-based (§4.1.3), using the `lsst_cells_v2` skymap. Coadds are built as weighted averages of the input warps. Each detector is assigned a weight that is the inverse of the mean of the variance reported in the Visit Summary dataset (§4.5.3). Following (Armstrong et al. 2024) for each cell in the coadd, a detector is excluded if its geometry does not overlap the cell completely. A detector is also excluded if more than 50% of its pixels are masked as artifacts, unless doing so would result in an empty cell, in which case this condition is ignored. Finally, pixels that did not get any contribution from the input visits are interpolated from neighboring pixels.

A key difference between cell-based coadds and coadds in DP1 is in the PSF representation. Single visit PSFs are evaluated at a point corresponding to the center of each cell and warped into a 35x35 pixel postage stamp, which ensures that at least 99.5% of the flux is captured. These postage stamps are then weighted by the same weights used for the coadd and averaged. The PSF within the cell is approximated as spatially constant. The pixels that have artifacts from any of the input visits have `INEXACT_PSF` flag set to denote that the PSF of those pixels may differ from that attached to the cell. The metadata about which visits contribute to the cell, what fraction of pixels are unmasked, PSF summary statistics are stored.

5.3.3. *Detection, Deblending, & Measurement*

Coadd detection and measurement follow a similar sequence of steps as in DP1, though numerous changes have been made to individual algorithms; the main change is that these algorithms now run on cell-based coadds, which have fewer PSF discontinuities within a cell. For this release, however, none of the measurement algorithms account for cell boundaries directly: any algorithm that requires a PSF model instead uses the model evaluated at the centroid of the object (and therefore of the cell containing that centroid) throughout. The processing proceeds as follows: peaks are first detected in single-band patch images, then merged across bands into a multiband peak catalog; this catalog is the input to deblending, which produces deblended (but not deconvolved) images for each object, and all subsequent object catalog measurements are performed on these deblended images.

The basic methodology for detection is largely unchanged from the description in Bosch et al. (2018), which we summarize here. To detect peaks, each image is convolved by a circular Gaussian approximation to the measured PSF to make a significance image. A footprint is a contiguous region of pixels above the detection threshold, which is set to detect isolated point sources above 5σ significance. Footprints can have multiple peaks, which are defined as pixels with values higher than all 8 neighbours (local maxima within a 3x3 box). Note that even for an isolated object, the significance of a detection is not precisely the same as the signal-to-noise from subsequent PSF model fluxes.

Pixels within footprints are marked with the `DETECTED` bit mask. The remaining pixels are used for background estimation and subtraction, excluding significantly negative-valued pixels which are likely to be residual artifacts. The algorithm also attempts to randomly place 100 "sky objects" within the noise-dominated pixels, reducing the number in crowded fields with low numbers of sky-dominated pixels if necessary. These objects do not appear in the final object catalog but are used to characterize the sky variance in intermediate steps.

Aihara et al. (2019) show that convolving the image with an approximation of the PSF reduces the image variance, so the pixel significance threshold must be adjusted accordingly to achieve the desired detection threshold. This factor is computed by measuring the standard deviation of the PSF model SNR of the sky objects.

Once detection has run in each available band, footprints and their peaks are merged in a fixed priority order, *irzygu*, which is roughly in order of expected SNR for typical objects. New peaks are only added if they are at least $1\farcs0$ further than an existing peak, and existing peaks are marked as detected in multiple bands if a new peak is found within $0\farcs3$; peaks within this range are ignored. After merging, probable spurious peaks detected in only one band are removed.

Deblending is performed by the `ScarletLite` python package. `ScarletLite` is an optimized implementation of the scarlet (Melchior et al. 2018) deblending algorithm with the assumption that the input images are aligned to the same pixel grid at a fixed resolution, as is the case in coadds. Prior to deblending, a new addition in DP2 is a per-band partial deconvolution of the input coadds to a common $\sigma = 0.8$ pixel integrated circular Gaussian PSF. This allows us to sub-divide large parent blends (footprints) into multiple smaller blends that are only connected by the wings of the PSF. These smaller sub-blends and the initialization of galaxy morphologies from the partially deconvolved images improves numerical convergence and performance, both in the goodness-of-fit of the final models and in memory usage and runtime. Users will notice that this has improved our coverage[39] in crowded fields from DP1; however, in the densest fields, an entire patch can still be a single blend and must be skipped due to excessive memory requirements.

Object measurements are conducted in three modes: independent per-band measurements, forced measurements in each band, and multiband measurements. Following Bosch et al. (2018), measurement begins by replacing all pixels within footprints with noise, allowing algorithms that need pixel values beyond footprint boundaries to measure without contamination from neighbors, whether or not those neighbors are part of the same blend. During measurement, isolated objects have their original pixel values inserted back into the image, while blended objects use the scarlet lite models as templates to re-distribute the original image values. This preserves the flux from the original image where objects have no overlap and dampens the noise elsewhere, since the flux is reapportioned to multiple objects.

The single-band measurement plugin system remains unchanged from DP1 and largely follows the description in Bosch et al. (2018). Measurements are made in every available band, and a reference band is defined from the same priority order as detection merging (*irzygu*) for a second round of forced measurements. These use the shape and position from the reference for consistency. Most algorithms have their free measurements reported in the final object table, but forced values are generally recommended.

[39] the fraction of each field with usable measurements

A variety of flux measurements are provided, including aperture fluxes, Kron (1980), and Gaussian-Aperture-and-PSF (GAaP Kuijken 2008; Kannawadi 2025) fluxes for consistent galaxy colors. Forward PSF model and Composite model (CModel) magnitudes (Abazajian et al. 2004; Bosch et al. 2018) are provided for each band.

MultiProFit (Taranu 2025) performs multiband forward modelling. Beginning in DP2, an exponential profile is fit first and used to initialize a subsequent Sérsic model (Sérsic 1963; Sersic 1968) fit, which improves convergence and reduces the prevalence of Sérsic fits with unrealistically large effective radii. Both exponential and Sérsic model fits have free centroids, which serve as alternative, multiband astrometry for objects. The Object table now includes uncertainties on the fitted ellipse parameters for both models (the x- and y-axis effective radii and correlation parameter ρ). These will be converted to uncertainties on more commonly used parameters (major/minor effective radii and axis ratios) for Data Release 1 (DR1).

Some additional derived parameters are included in the object catalog, mainly for object classification. "extendedness" is a binary classifier, set to 1 if the PSF to CModel flux ratio is below 0.985. "sizeExtendedness" is a quasi-probabilistic morphological classifier reported as $1 - \exp(-0.5 * ((T_{\rm HSM} - T_{\rm PSF})/T_{\rm PSF}^{0.5})^2)$, where $T_{\rm HSM}$ and $T_{\rm PSF}$ are the sum of the squares of the x- and y-axis moments for the object from the HSM algorithm and for the PSF model, respectively. As an alternative, DP2 adds `model_extendedness`, which is a continuous, heuristic classifier that uses both the exponential model effective radius and the ratio of the Sérsic model to PSF model flux. It is available for each band and separately for the sum of the *griz*-band fluxes.

As for DP1, photometric redshifts are not computed or included in the object table but are planned as a supplementary user-generated product; further details on the performance of these algorithms are given in §6.

5.3.4. *Metadetection*

The ShearObject catalog using the multiband metadetection algorithm (Sheldon et al. 2020) is new to DP2. It creates a separate catalog of detections that do not correspond to entries in the Object catalog. Each 150x150 cell in *griz* bands is extracted with a 50 pixel border to make 250x250 pixel images, that are then passed to metadetect. Internally, metadetect makes counterfactual images by deconvolving by the original PSF, shearing the image and reconvolving to a new PSF using the `fitgauss` reconvolution algorithm. Metadetect coadds *riz* bands to form a coadd for detection, and measures shear using the `gauss` algorithm by fitting a 2D elliptical Gaussian to *r*, *i*, *z* coadds. Additionally, for these detections, it also reports force-photometered fluxes in *g*, *r*, *i* and *z* bands for these detections measured using the `pgauss` algorithm. There is no deblending performed prior to these measurements.

5.4. *Variable & Transient Source Processing*

5.4.1. *Difference Image Analysis (DIA)*

DP2 difference imaging analysis did not change significantly relative to DP1. DP2 DIA uses the decorrelated Alard and Lupton image differencing algorithm (Alard & Lupton 1998; Reiss & Lupton 2016). Both positive and negative `DIASource`s are detected in the difference image at 5σ. We fit a dipole model to sources containing both positive and negative peaks, and infer trail parameters from the second moments of the PSF. After rejecting likely artifacts (§5.4.2), we associate `DIASource`s into `DIAObject`s with a 1.″0 matching radius.

5.4.2. *Artifact Rejection*

DIA analysis is prone to false positives. We employ several layers of artifact rejection, although we err on the side of completeness rather than purity in Data Release Processing.

Artifact rejection in DP1 consisted of removing sources with pixel flags artifacts, non-astrophysical trail lengths, and unphysically negative direct fluxes. In DP2 we also reject sources that land on the modeled positions of diffraction spikes and apply morphological cosmic ray rejection on both the direct and difference images. Further, we flag but do not remove occurrences of five or more `DIASources` that occur in a straight line with `glint_trail`, indicating that they may be due to rotating orbital debris (Karpov & Peloton 2022).

We developed a machine-learned reliability score to assess whether `DIASources` were astrophysical. This model is a convolutional neural network that was trained on template, science, and difference-image postage stamps around `DIASources` using a combination of simulated fake injections and human-labeled detections. (Acero-Cuellar et al. 2026) details the model architecture, training, and performance of the real-bogus model. The v0.1 model used in DP1 was again used during DP2 processing. No cuts were applied on reliability for the DP2 `DIASource` catalogs, so users interested in higher purity should apply a minimum reliability threshold.

5.4.3. *Light Curves*

Light curves are compiled from multiepoch forced photometry performed on both the Visit Images and the difference images. The forced photometry per-

formed on the difference-images (the `psfDiffFlux` column of the ForcedSource table) measures only the time-variable component of the flux without contamination from static nearby sources. Forced photometry carried out directly on the Visit Images can pick up flux from blended or nearby static objects, and the severity of this contamination depends on the seeing. The centroids of objects with multi-epoch forced photometry are drawn either from coadd-based object positions or from the DIAObjects, which are constructed from the DIASources associated from detections in the difference images.

5.5. *Solar System*

5.5.1. *Source Association*

The solar system association input was a 2025 March 13 snapshot of the MPC Orbit database[40], limited to objects whose observational arcs exceeded two days. Associations were made independently for each detector and visit using positions predicted by Sorcha (Merritt et al. 2025; Holman et al. 2025) at the visit epoch from this orbit catalog. Predicted objects were first restricted to the detector footprint. Measured DIA Source (§4.2.2) coordinates were represented as unit vectors on the celestial sphere and stored in a k-d tree. For each predicted position, the algorithm identified the nearest detection within $1''$. Candidate pairs were considered in increasing separation, accepting a pair only if neither member had already been assigned. Thus, this is a positional, one-to-one match; it uses neither brightness, color, morphology, nor rate and does not calculate a probabilistic association score. The delivered `diaDistanceRank` is set to one for every row and should not be used as a measure of match ambiguity. Chance associations remain possible because the match does not use brightness or detection reliability, and are more likely where the DIA source density is high. Users constructing a high-purity sample should consider detection reliability, angular and cross-track residuals, consistency with the predicted magnitude, and the uncertainty of the input orbit. Detection-level reliability information requires reference to the DIA Source catalog (§4.2.2).

5.5.2. *Photometry and Phase Curves*

The solar system Object photometric fields were derived independently in each of the *ugrizy* bands from PSF fluxes in the associated detections. For a flux f in nJy, the conversion used was $m_{\rm AB} = 31.4 - 2.5\log_{10} f$, with magnitude uncertainty $1.085736\,\sigma_f/f$. Non-positive fluxes therefore yield non-finite magnitudes and were removed before fitting. Apparent magnitudes were reduced to unit heliocentric and topocentric distance using $m(1,1,\alpha) = m - 5\log_{10}(r\Delta)$, where r and Δ are in au.

The reduced magnitudes were described with the H, G_{12} phase function of Muinonen et al. (2010). The DP2 processing fixed $G_{12} = 0.5$ and fit only the absolute magnitude H, since the short observing interval and limited phase-angle coverage do not generally support a stable two-parameter fit. Accordingly, valid per-band `G12` values are all 0.5, while `G12Err` and the H–G_{12} covariance are undefined. A 0.05 mag uncertainty floor was added in quadrature to the measurement uncertainties to reduce sensitivity to unmodeled scatter, including rotational variability. After removing non-finite magnitudes, an initial fit with a robust loss was used to reject residuals beyond 10σ, followed by a conventional least-squares fit to the retained observations. The per-band `nObsUsed` field records the retained count and should be compared with `nObs` when assessing clipping or invalid photometry.

Valid H estimates were obtained for 35,268, 196,015, 224,888, 231,890, 203,341, and 90,789 objects in *ugrizy*, respectively. Among objects with sufficient observations for a fit, failures range from 53 in y to 1,034 in g, or less than approximately 0.5% in every band. Depending on band, 1–5% of the eligible measurements were removed as invalid or clipped. There are 25,706 objects with no fitted H in any band; many are represented by only one detection and hence have insufficient information for a fit. The median fitted H values are 18.7, 18.2, 17.5, 17.3, 17.2, and 16.6 mag in *ugrizy*. These statistics are dominated by main-belt objects and are also shaped by the DP2 footprint, cadence, phase-angle coverage, and band-dependent detection selection; they should not be interpreted as intrinsic population size distributions.

6. PERFORMANCE CHARACTERIZATION AND KNOWN ISSUES

In this section, we present a characterization of the performance of the DP2 data products, and highlight known issues, as understood at the time of publication. The analyses presented here provide an overview of the data quality rather than a comprehensive validation; where appropriate, more detailed analyses are deferred to dedicated future papers. Further information on the performance and known issues of DP2 can be found in the Data Quality, Image Artifacts, and Known Issues sections of the Release Documentation at

[40] Small bodies discovered after 2025 March 13—including those found in Rubin data—are therefore absent from the solar system Source catalog §4.2.9

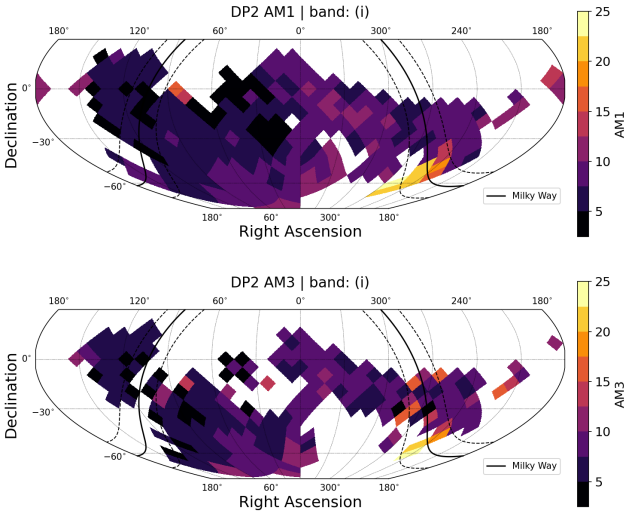


Figure 7. Whole sky plots showing the per-healpix median of the standard deviation in the separation across all *i*-band visits between objects at a given distance from each other, which are called the "AMx" metrics. Lower values indicate a high degree of astrometric consistency between visits. The top panel shows the values for objects separated by 5 arcmin, while the bottom shows the values for objects at 200 arcmin, demonstrating that in many areas of the sky we are achieving the 10 mas design goal for AM1 and the 10 mas *stretch* goal for AM3. The difference in sky coverage is due to the fact that some regions of the sky do not have a sufficient number of objects separated by 200 arcmin.

https://dp2.lsst.io, which will continue to be updated as our understanding of the data and its nuances improves.

We are awaiting the subsection on sensor anomolies.

6.1. *Astrometry*

Astrometric performance is characterized by measuring internal consistency as well as accuracy with respect to an external reference, in this case the Gaia DR3 catalog. Metrics are calculated using isolated stars with magnitude between 17 and 21.5. For this set of objects, we find that the median offset from Gaia is within 2 mas in about 85% of tracts and within 5 mas for about 99% of tracts. Any offset within this margin is believed to be caused by color-dependent selection effects in our metric calculation, and is expected to be lower once differential chromatic refraction is included in the astrometric model.

Internal consistency was improved significantly relative to DP1 by including the Gaussian processes turbulence correction. The degree of internal consistency can be seen by looking at the "AMx" metrics, which show the median over a field of the standard deviation in the separation across all visits between objects at a given distance. Figure 7 shows this metric for objects at 5 arcmin separations and at 200 arcmin separations. The top panel show that DP2 is reaching the design goal of repeatability below 10 mas in many regions of the sky. Furthermore, we are achieving almost the same repeatability at 200 arcmin separations, for which 10 mas is actually the stretch goal.

6.2. *Stellar Photometry*

Comparison with an external photometric catalog is a key test of photometric quality, particularly of the absolute flux scale. Such a comparison is, however, far from straightforward, not least because any external catalog will also have its own calibration systematics. Further, outside the deep fields, any comparison would be restricted to relatively bright stars and would therefore not be representative of the full DP2 depth. A proper treatment of these limitations is beyond the scope of this paper, and is left to a dedicated future paper on photometric performance. Instead, we assess the internal consistency of the DP2 photometry using the repeatability of per-visit measurements and the stellar loci measured on the Deep Coadds. Both are differential measures (i.e., between visits or between bands), so they test the relative photometric calibration but are insensitive to any error in the absolute flux scale.

Since most of the DP2 footprint was observed multiple times, the repeatability of photometric measurements of the same star across visits provides a test of the consistency, and hence reliability, of the per-visit photometric calibration. We measure this repeatability using the PSF photometry of bright (SNR $>$ 200), isolated, unresolved objects, the vast majority of which are stars. To keep the calculation computationally feasible, it is performed per tract: for each star with at least three measurements from the visits covering a given tract, we calculate the standard deviation of those measurements, and then take the median over all stars in the tract. This corresponds to the PA1 metric defined in the LSST Science Requirements Document (Ivezić & The LSST Science Collaboration 2018). Following recalibration using the FGCM fits, the median PA1 values across all tracts with coverage are 9.1, 7.1, 6.8, 5.9, 5.8, and 6.2 mmag in the *ugrizy* bands, respectively, meeting the minimum specifications of the Science Requirements Document.

Figure 8 shows the *ugriz* stellar loci of unresolved objects in the Deep Coadds, selected using the Object table's `extendedness` parameter (§5.3.3). The tightness of these loci indicates the good quality of the band-to-band relative photometry in the Deep Coadds.

6.3. *Differential Chromatic Refraction*

DCR is the wavelength-dependent astrometric displacement arising from atmospheric refraction. The dis-

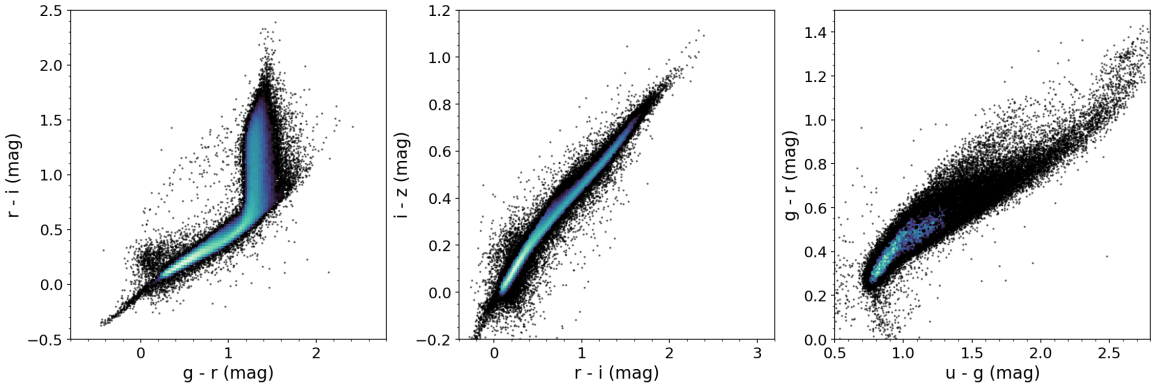


Figure 8. Stellar loci measured on the Deep Coadds using PSF magnitudes from the DP2 Object catalog. From left to right, the *gri*, *riz*, and *ugr* loci are shown for random subsets of bright (i-band SNR > 200; $grzy < 22.5$ mag), isolated, unresolved objects at Galactic latitude $|b| > 20°$, the latter to limit the effects of dust reddening; the *ugr* panel additionally requires u-band SNR > 50. The first two panels each contain 253,926 objects, and the third contains 41,968. The color scale is logarithmic, with lighter colors indicating a higher density of objects, and is set independently for each panel; bins containing fewer than twenty objects are shown as individual points.

placement is directed along the parallactic angle and increases toward shorter wavelengths, producing a color-dependent angular offset the size of which differs between bands.

The presence of DCR in LSSTComCam data was previously discussed in the DP1 release paper (Vera C. Rubin Observatory Team et al. 2026). Here, we characterize the effect in LSSTCam data using two representative sky regions: the Extended Chandra Deep Field-South (ECDFS) centered at (R.A., decl.) = (53.16, -28.10) and a field near the ecliptic centred at (R.A., decl.) = (320.0, -15.0), respectively. For each region, we consider all Deep Coadd images from tracts within 5 degrees of the field center.

The extent of any displacement due to DCR is dependent on airmass, meaning that for a patch it changes at the per-visit level. For each visit, we calculate the angular displacement of sources relative to their positions their positions in the matched reference catalog and projected onto the directions parallel and perpendicular to the parallactic angle. We examine both components as a function of source $g - i$ color; the perpendicular component serves as a control for effects not aligned with the expected DCR direction.

When displaying the resulting DCR effect, the coadd with the highest DCR metric is selected for each field. In the case of the ECDFS, that specific coadd contains 7 visit images, spanning an airmass range of 1.02-1.36 with a mean of 1.14. The corresponding ecliptic coadd contains 8 visit images spanning 1.05-1.70 with a mean airmass of 1.26. As shown in Figure 9, both fields exhibit a color-dependent displacement parallel to the parallactic angle, with no corresponding trend in the perpendicular component. The color dependence rules out a

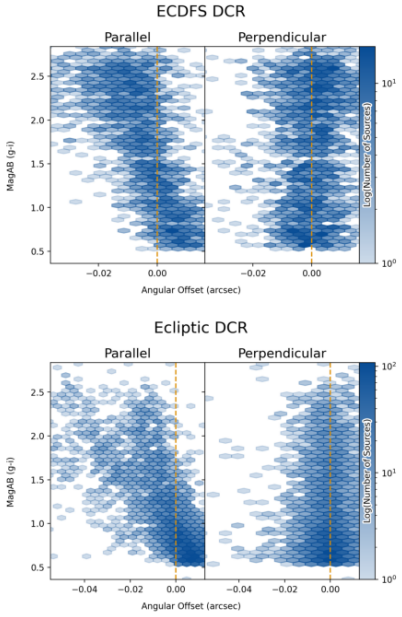


Figure 9. Differential Chromatic Refraction (DCR) measurements for the ECDFS (top) and ecliptic fields (bottom) in LSSTCam data. For each field, the angular displacement of sources is measured relative to their positions in the matched reference catalog and projected onto directions parallel and perpendicular to the parallactic angle. The binned source density (see colorbar) is shown as a function of $g - i$ color and angular displacement, revealing a color-dependent correlation along the parallel direction. This, and the lack of a corresponding trend along the perpendicular direction, indicates the presence of DCR in the LSSTCam data. While we consider all Deep Coadds within five degrees of each field center for our analysis, for display purposes we only plot the DCR values from all matched sources in the coadd with the highest overall DCR metric. **Note:** We are aware the axis labels etc are too small. This will be addressed.

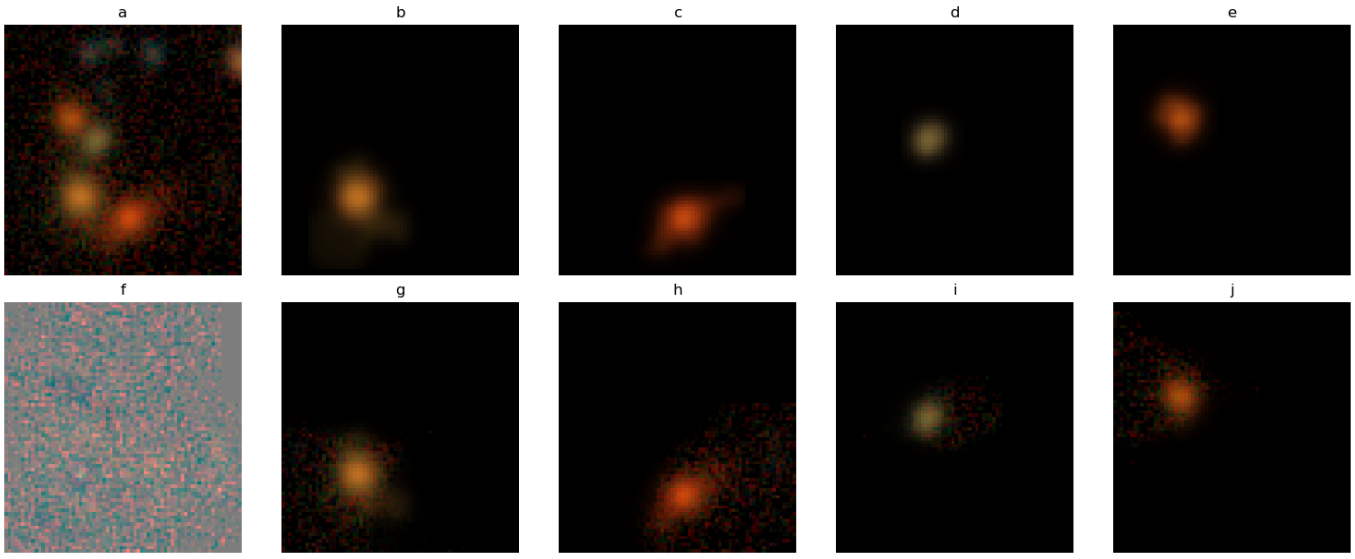


Figure 10. Example multi-object blend. The top row shows the observed multiband image (panel a) followed by the scarlet lite models of the four brightest sources (panels b through e). The bottom row shows the residual image when the scarlet lite models are subtracted from the observation (panel f) followed by the flux-conserved models (panels g through j), created by using the scarlet lite models as templates to re-distribute flux from the observation so that the total flux is conserved. The flux-conserved models, injected in a noise-replaced image, are what the science pipelines use for measurement.

simple camera-dependent offset and is consistent with the DCR signature previously identified in LSSTComCam. This effect is being addressed through the development of DCR-specific templates and coadds designed to mitigate the observed offset in future data products.

6.3.1. *PSF accuracy*

A full analysis of the PSF characterization of the Rubin Observatory data is currently being written (Léget 2026), so here we present only the key findings. One of the key findings is that the PSF shape correlates with variations in height across the focal plane (see Laliotis et al. 2026 for more details). This explains why ITL sensors, which exhibit a curved shape, are more difficult to model than E2V's flat sensors, and why a fourth-order model is justified, as noted previously in the DP1 paper. As described in Pai et al. (2026), the LSSTCam PSF is chromatic (i.e., wavelength-dependent) and appears to be compatible with DCR. As shown in Pai et al. (2026), we are able to model it in the same spirit as the DES Y6 PSF analysis, but this was not enabled during DP2 processing because the connections to the downstream processes were not in place. In addition, DP2 has an order of magnitude more visits than previous processing campaigns, revealing a new variety of PSF-related issues, including annealing patterns in E2V sensors — visible in the u- and *g*-band stacks — and fringing in the z- and y-bands. All are expected to arise from background subtraction. Moreover, PSF residuals in DP2 show structure that correlates with stellar density, with the largest biases appearing in the plane of the Milky Way, the LMC, and the SMC. This stellar-density dependence also explains the observed differences in PSF residuals between bands, since more crowded bands show larger residuals. The underlying cause is again believed to be related to background subtraction in dense fields.

6.4. *Deblending*

As mentioned in §5.3.3, a new feature of DP2 processing is the introduction of a partial deconvolution step prior to deblending. Including this step significantly decreases the number of blends. As a result, in extragalactic fields such as COSMOS, just over half of the objects are not blended in the partially deconvolved image, allowing the deblender to quickly and efficiently recover the flux for objects isolated in observed image and partially deconvolved image. Most of the remaining blends are small, with only a handful of objects, which scarlet lite deblends fairly easily (see Figure 10). Approximately 10% of objects come from blends with more than 5 children.

The biggest failure modes are resolved grand-design spirals, which are shredded into multiple poorly modeled objects, and undetected objects in the wings of brighter objects. The deblender's primary limitation is its memory usage, which in crowded fields can become large, and for DP2 forced us to skip deblending near the galactic plane and a few other dense regions. We are currently working on algorithms to allow us to process crowded fields and expect most fields to be successfully

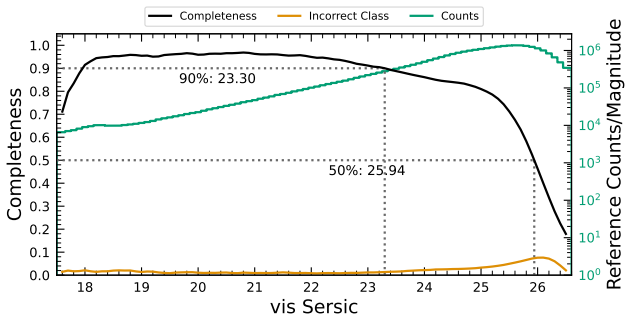


Figure 11. Completeness and incorrect classification fraction as a function of Euclid VIS-band Sérsic magnitude. The "Incorrect Class" line shows objects that are matched but classified incorrectly by `griz_model_extendedness`.

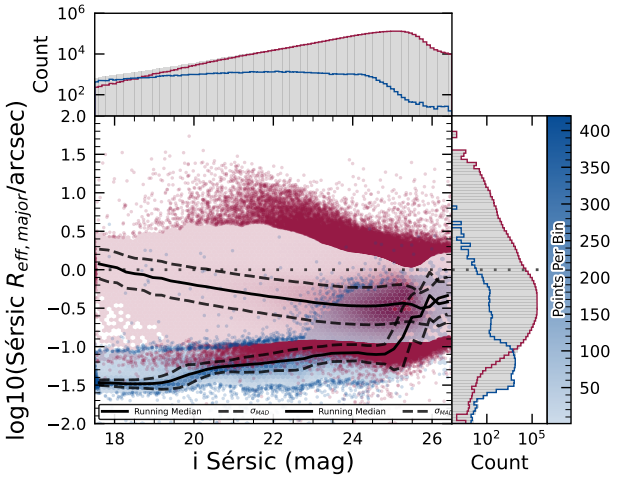


Figure 12. Sérsic major-axis effective radii as a function of Sérsic i-band magnitude for the ECDFS field. Red points are classified galaxies with `griz_model_extendedness` > 0.2, and the hex binned density transitions to a scatter plot at the 0.5 and 99.5 percentiles in size.

processed in future releases. For users wanting a clean sample of objects that are not blended in the partially deconvolved image, meaning objects that are either fully isolated or that overlap neighbors only in the observed image, we recommend selecting objects with `deblend_blendNChild` equal to zero or with `detect_isIsolated` set.

6.5. *Coadd Depth, Completeness and Purity*

The depth of coadds can be estimated in several different ways, as well as separately for point and extended objects. One method is to simply measure statistics like the median/mean magnitude of objects in a narrow range of significance, i.e. at $5\sigma \pm 0.1\sigma$. One drawback of this method is that uncertainties in coadd fluxes are typically underestimated, and by a larger factor for extended objects. Nonetheless, it is fairly straightforward to implement this method and correct the error estimation, at least for point sources.

Injecting synthetic objects into real data is another method that was used for DP1. However, the switch to cell-based coadds necessitated changes to the injection methodology for DP2 and a large-scale injection campaign has not yet been completed.

The third method for measuring completeness is to match objects against comparable or deeper reference catalogs, where higher-resolution data from space-based telescopes is preferred. For this method, as for DP1, we use a greedy, probabilistic matching algorithm that matches reference objects in order of descending brightness to the most probable target within a $0\farcs5$ radius.

Figure 11 shows completeness in the ECDFS field compared to Euclid Q1 (Euclid Collaboration et al. 2026a,b), along with the proportion of objects whose star/galaxy classification disagrees with Euclid. Euclid stars are defined solely with a `mumax_minus_mag < -2.6` cut as recommended by Euclid Collaboration et al. (2026c), whereas DP2 stars are selected with `griz_model_extendedness` < 0.2. Completeness is generally high, peaking above 96.5% for $20 < \mathrm{mag_{VIS}} < 21$, with a misclassification rate just under 1%. Purity, on the other hand, peaks at bright magnitudes but dips below 93.3% at the same magnitude range.

Completeness and purity figures are somewhat dependent on the choice of flags to exclude possible spurious detections. For Euclid, we select only objects with `spurious_prob`< 0.2. This cut excludes more objects than the default value of 0.5 recommended by Euclid Collaboration et al. (2026b). Visual inspection of the coordinates of Euclid detections with $0.2 <$ `spurious_prob` < 0.5 revealed that most are clustered around bright stars and their diffraction spikes.

Similar problematic detections are found in the DP2 catalogs, many but not all of which are excluded using `coord_flag`. Not all of these flagged objects are necessarily spurious detections, but as with Euclid, many cluster around bright stars and have overestimated fluxes due to their footprints including part of the extended wings of the star's PSF.

6.6. *Galaxy Photometry and Sizes*

One of the main improvements in DP2 galaxy measurements is the use of an exponential fit to initialize the multiband Sérsic fits. Figure 12 shows the distribution of Sérsic size as a function of Sérsic i-band magnitude in ECDFS. While there are still some overly large galaxy outliers, the proportion is fairly small. The 99.5th percentile size increases slightly from 18th to 22nd magnitude, where the outliers are most problematic, but falls for fainter objects. This is likely partly due to improved

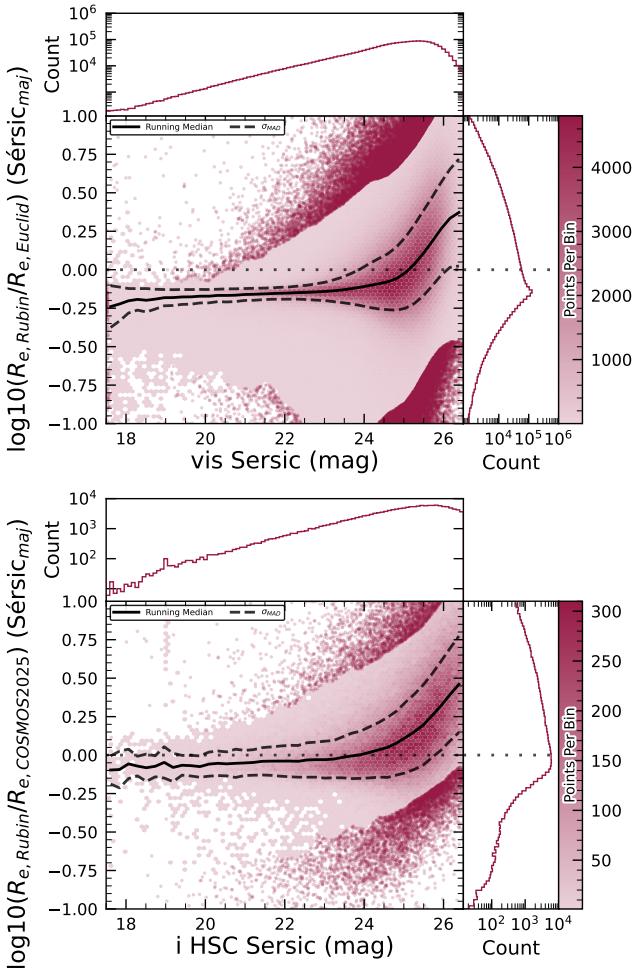


Figure 13. Sérsic size ratios for galaxies in ECDFS and COSMOS, as measured and classified in Euclid Q1 and COSMOS2025, respectively.

deblending as well; the Gaussian model in ShearObject, which does limited masking of neighbors rather than multi-object modelling, has a considerably larger fraction of "superspreaders".

Not sure if we want to be more quantitative about this; I can produce figures but it may be TMI.

Using `griz_model_extendedness`, the Sérsic sizes of stars are also consistently smaller than galaxies until 25th magnitude. At 24th magnitude, however, the number counts of stars drop precipitously, as the classifier's completeness drops to maintain purity. The cause of the significantly nonzero sizes of stars is not entirely understood, but the median sizes are consistent with other forward modelling algorithms, including from CModel's exponential fits (which use an independent double shapelet PSF model) and from the Gaussian model in ShearObject.

Figure 13 shows the ratio of Sérsic sizes measured in DP2 with two space-based catalogs - Euclid Q1 in ECDFS and COSMOS2025 (Shuntov et al. 2025) in COSMOS. In both cases, the median size ratio remains fairly consistent to about 24th magnitude (filter dependent), and the scatter is remarkably small for ECDFS/Euclid especially. There is a shallow slope in the median such that bright galaxies have underestimated sizes. This is likely due to a combination of reasons, most notably the known background oversubtraction which is less likely to be an issue in space-based data. Faint galaxies have substantially overestimated sizes, unsurprisingly, since by $i = 25$ median effective radii (from COSMOS2025) drop to about $0''.22$ - barely larger than the LSSTCam pixel size.

One as-yet unexplained feature is the systematic offset between DP2 and Euclid Q1 sizes, which is also seen in EDFS. The Euclid VIS band covers most of the same wavelength range as the *griz* filters combined, and so this seems unlikely to be solely due to a mismatch in wavelength coverage. By contrast, the Sérsic parameters measured in COSMOS2025 are from four JWST NIRCAM filters (F115W, F150W, F277W, F444W), only the bluest of which even overlaps partially with the *y* band.

6.7. *Isolated Star Stellar Motions*

Due to the short time span of the DP2 observations, the precision of the proper motion and parallax measurements is highly driven by the external reference catalog, Gaia DR3, which is included in the stellar motions calculation where possible. The majority of isolated stars with magnitudes below about 21 are in the Gaia catalog: for these objects, the proper motion uncertainty is on the scale of a few mas. For fainter objects for which Gaia is not available, on the other hand, even well-sampled objects have proper motion uncertainties on the order of $10 - 20$ mas, going higher as signal to noise, sampling frequency, and the time span sampled go down. Regardless, for the majority of tracts, the time spanned between the first and last visits included in DP2 is less than 100 days, meaning that proper motion and parallax measurements should be treated as a guideline for what will be available in future Data Releases, rather than a starting point for serious analyses.

6.8. *Solar System*

There are 299,343 ssObjects (i.e., objects measured at least once by Rubin) in DP2, representing almost 20% of the known solar system. Although 136,608 objects are observed at least once in each *griz* band, detailed analysis of asteroid colors would be premature with DP2.

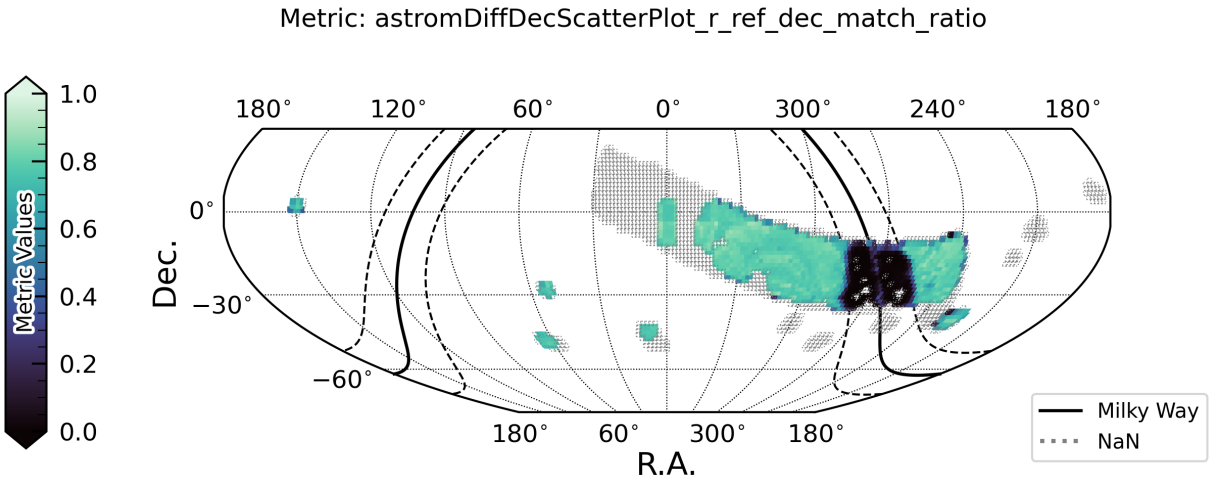


Figure 14. Reference-source match ratio across the DP2 survey footprint. For each tract/patch, the color scale denotes the fraction of reference sources successfully matched to LSST detections. Regions with no matched sources are shown in black, and the Galactic plane is indicated by the solid black curve. Low match ratios are concentrated along the Galactic plane. **Note:** There is a known bug that leads to the empty tracts in this figure; this will be addressed prior to submission to the journal.

Observations with different filters are not simultaneous, meaning the objects' rotational flux variations add scatter to the measurement of their colors. Given sufficient data, this can be overcome by fitting a rotational light curve for each individual object. However, most objects in DP2 have far too few measurements for a light-curve fit. Hence, individual colors and H magnitudes are uncertain at the level of rotational flux variations, which are typically much larger than the nominal uncertainty `HErr`, since the latter accounts only for the errors on instantaneous flux measurements.

The ensemble colors of the ssObjects are broadly plausible for an asteroid sample: median $g - r \sim 0.58$, $r - i \sim 0.18$, and $i - z \sim -0.01$ mag. Implausible outliers are rare ($\sim 0.25\%$ of all color measurements) and likely result from sparse phase fits heavily affected by individual photometric outliers. These photometric outliers may result from self-subtraction when the template images are shallow and contain flux from the asteroid. Like the problem of insufficient measurements for rotational light-curve solutions, this issue of shallow templates causing self-subtraction and occasional photometric outliers should naturally disappear over future Data Releases based on larger numbers of input images.

There are 8.1 million diaSources associated with solar system objects in DP2, representing about 2% of all MPC measurements. There is a persistent ~ 30 mas bias between DP2 astrometric measurements and ephemeris positions of MPC orbits. When Rubin astrometry was compared instead to orbits fit only to Gaia astrometry, the ~ 30 mas bias disappeared. After substantial investigation and discussion with both the MPC and JPL, we concluded that the bias is **not** in Rubin astrometry (to ≤ 1 mas), but in the orbit catalog. While 30 mas is small, the median astrometric uncertainty on a DP2 diaSource associated with an ssSource is just 56 mas, making 30 mas a substantial offset in the context of Rubin precision, thus highlighting the extent to which Rubin astrometry can improve asteroid orbits over time.

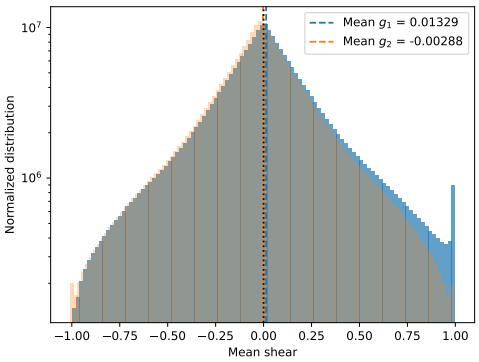


Figure 15. Distribution of `gauss_g1` (blue shaded histogram) and `gauss_g2` (orange shaded histogram) values for the ShearObject catalog's no-shear subcatalog; overlapping histograms appear as gray shading. The mean values of `gauss_g1` and `gauss_g2` are marked as colored vertical dashed lines (see legend); the black dashed line indicates zero. Over large areas such as those covered by DP2, the mean shear should be zero. While small, the mean shears are significantly nonzero. We currently attribute this offset to objects in the tails of the distribution, which are typically associated with bright stars and nearby galaxies.

6.9. *Crowded Fields*

Crowded fields represent a key challenge for the LSST Science Pipelines, pushing the current limits of source detection, deblending, and photometry. They are defined as regions in which the stellar source density exceeds the resolving capability of current algorithms. In such regions, the pipeline produces large, contiguous detection footprints encompassing multiple astrophysical sources rather than assigning distinct footprints to individual objects; the corresponding pixels are flagged as `NOT_DEBLENDED`. We therefore use the proportion of pixels flagged as `NOT_DEBLENDED` as a proxy for the area of the survey that is associated with crowded fields.

Overall, out of the 2.2 trillion pixels comprising the total DP2 sky area, 12.9% are associated with crowded fields via the aforementioned definition. This fraction, however, obscures the highly localized nature of crowded fields within the survey. Specifically, 1,750 of 2,174 total tracts (i.e., 80.5%) have less than 1% of their pixels flagged as `NOT_DEBLENDED`. Of the remaining 424 (i.e., 19.5%) tracts , 90% are found within 14.5 degrees of the Galactic plane where the stellar densities are the highest. The concentration of processing failures in these regions is further illustrated by the fraction of reference sources successfully matched to LSST detections within each tract and patch. Figure 14 shows this match ratio across the survey footprint. The lowest match ratios, including tracts and patches with no successful matches, occur almost entirely along the Galactic plane, indicating that these unmatched sources are associated with failures of the standard LSST processing in these densely populated fields. Away from the Galactic plane, match ratios remain uniformly high. Thus, although crowded fields constitute a substantial fraction of the total flagged pixels, the associated processing challenges are concentrated predominantly in a relatively small region of the survey footprint near the Galactic plane.

As the survey continues, crowded fields will comprise an increasingly large fraction of the observed sky. Robust processing of these fields is therefore essential to fully realizing the scientific goals of LSST. To address this challenge, a dedicated crowded-field processing pipeline is currently under development. Early results are promising: regions that appear as merged detections in the standard pipeline can instead be successfully deblended into their constituent sources, enabling individual detection and measurement. The pipeline is currently undergoing testing, with integration into the LSST Science Pipelines for general use targeted within the next year.

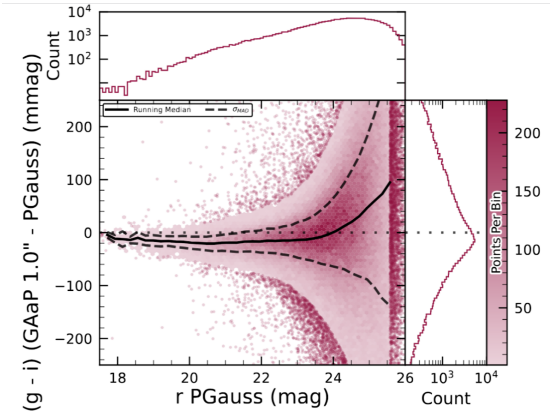


Figure 16. Comparison of $g - i$ colors measured with the `pgauss` algorithm in the no-shear ShearObject catalog and with 1.″0 GAaP photometry in the matched Object catalog, as a function of r-band `pgauss` magnitude. The main panel shows the color difference (GAaP − `pgauss`) in mmag, with the density of objects shown as shaded hexagonal bins and individual objects shown as points in sparsely populated regions. The solid black line shows the running median and the dashed lines show the running $\sigma_{\rm MAD}$; the dotted line marks zero difference. The top and right panels show the distributions of r-band magnitude and color difference, respectively.

6.10. *Shear Object catalog*

As described in §4.2.8 and §5.3.4, the ShearObject catalog is produced by running the metadetection algorithm on the coadded images. It contains the lensing shear measured as ellipticity on the no-shear counterfactual, plus four subcatalogs from the counterfactuals in which a linear shear has been artificially applied; the latter are used to compute the shear response and selection biases. A full validation of the Shear Object catalog is beyond the scope of this paper and will be the focus of follow-up studies. Nevertheless, a basic check is possible by exploiting the expectation that over large areas such as those covered by DP2, the mean shear of the no-shear counterfactual should be zero. We therefore measure the mean `gauss_g1` of detections in the no-shear counterfactual, after excluding any with a flag set or with `is_primary=False`. In doing so, we obtain a mean `gauss_g1` value that, while small, is significantly nonzero. This offset persists when we restrict to the best detections (`mfrac<0.1` and `gauss_snr>10`) and when we apply inverse-variance weights to downweight fainter detections. Following further investigation, we attribute this offset to objects in the tails of the `gauss_g1` distribution, which, from visual inspection, are typically bright stars and nearby galaxies. Such objects may therefore need to be masked out by hand prior to any shear analysis.

Further validation of the Shear Object catalog can be performed by matching the no-shear catalog to the Object catalog. We stress that this is suitable for validation only; matching reintroduces the selection biases that metadetection is designed to remove, so the Shear Object subcatalogs must not be matched to one another or to the Object catalog in any shear analysis. For this check, we compare the `pgauss` $g - i$ color of the no-shear catalog to the GAaP $g - i$ colors in the Object catalog. As these are designed to measure comparable quantities, we expect the two to largely agree, so any major disagreement would point to a problem with one of the two measurements. This comparison, shown in Figure 16, reveals good agreement for $r \lesssim 23$, with a median offset (GAaP − `pgauss`) of ~ -20 mmag; at fainter magnitudes, both the offset and the scatter increase.

6.11. *Difference Imaging*

DP2 includes over 12,000 visits with `DIASources`. As discussed in 4.2.2, deblending is not yet implemented. Per 5.4.2, reliability scores are calculated for each of these, but no cuts are applied. As a result, the majority of `DIASources` have low reliability scores (< 0.5) and are not astrophysical in origin. Common causes of low-reliability `DIASources` include low-quality Template Coadds (e.g., single visit self-subtractions) and poorly constrained background sky levels due to scattered light or crowded fields.

In some circumstances, spurious `DIASources` can appear in a grid pattern reminiscent of the gaps between detectors on the focal plane. This can occur when Template Coadds have a sudden change in PSF size. This discontinuity cannot be well represented with a smooth matching kernel during difference imaging source measurement, and a too-large PSF can produce spurious positive flux `DIASources` in the narrow grid strips where the PSF is small. There is presently no flag or column to identify these, but we expect it to be less prevalent for Template Coadds that are deeper (i.e., those with larger numbers of overlapping input visits).

Most of the visits with `DIASources` in DP2 have fewer than 2000 `DIASources` with reliability > 0.5. The outliers with excess "reliable" `DIASources` include: (1) a handful of visits listed in the Release Documentation Known Issues when the telescope was moving while the shutter was open, (2) visits affected by the template PSF discontinuity problem described above, and (3) visits in crowded and/or nebulous fields where it is challenging to robustly measure the background flux level.

6.11.1. *DIA Source detection completeness*

Following the approach in DP1 (Vera C. Rubin Observatory Team 2026), we assess the performance of difference imaging by injecting synthetic sources into a subset of DP2 science images, carrying out difference imaging, and measuring properties of the recovered detected synthetic sources. We consider a selection of i-band images from the deep ECDFS field for this investigation. We inject both "hosted" and "hostless" synthetic sources, i.e., point sources embedded in a host galaxy and isolated point sources, each with randomly sampled size and brightness. The hosted sources additionally have randomly selected positions within their host galaxy and position angles.

Bruno and Meredith are working on this fake injection campaign and will write up results soon. We should be able to generate 2–3 figures similar to DP1 (Vera C. Rubin Observatory Team 2026) Figures 37, 38, and 39.

7. RUBIN SCIENCE PLATFORM

The RSP (Jurić et al. 2019) is a powerful, cloud-based environment for scientific research and analysis of petascale astronomical survey data and serves as the primary interface for scientists to access, visualize, and conduct next-to-the-data analysis of Rubin and LSST data. Built around a "bring the compute to the data" principle, it eliminates the need for users to download massive datasets, instead providing Rubin data rights holders with secure, single-sign-on access to data and services via three distinct user-facing Aspects: a *Portal*, which facilitate interactive exploration of the data; a JupyterLab-based *Notebook* environment for data analysis using Python; and an extensive set of *API*s that enable programmatic access to both data and services.

The three Aspects are designed to be fully integrated, enabling seamless workflows across the RSP. The RSP is supported by a number of backend services, including databases, files, and batch computing. Support for collaborative work through shared workspaces is also included in the RSP. The data products described in §4 are accessible via all three Aspects, and the system facilitates operations such as starting a query in one Aspect and retrieving its results in another. Figure 17 shows the RSP landing page in the Google cloud.

A preview of the RSP was first launched on Google Cloud in 2022, operating under a shared-risk model to support DP0 (O'Mullane et al. 2024b). DP1 built on that early preview to bring major upgrades including the hybrid on-premises USDF/cloud deployment, a broader set of International Virtual Observatory Al-

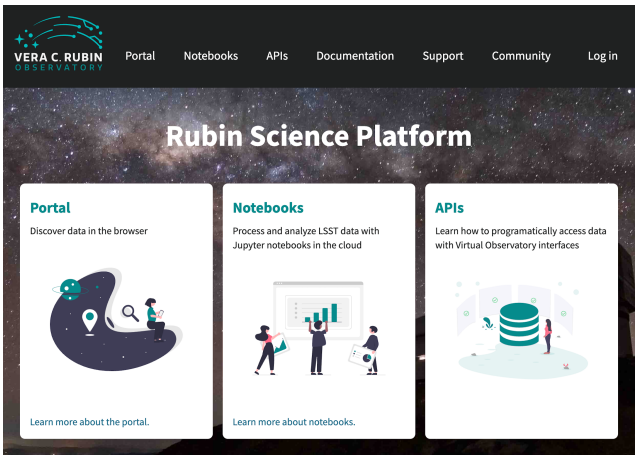


Figure 17. The RSP landing page at https://data.lsst.cloud/ showing the three user-facing Aspects as well as links to documentation and support information.

liance (IVOA) API services (SIAv2, SODA cutouts, HiPS, WebDAV, alongside the existing TAP), and a fully integrated, read-only Data Butler (S3 store + AlloyDB registry) across all three RSP Aspects.

The size (3.5 TB) and complexity of DP1 did not, on their own, necessitate the full functionality of the RSP for analysis; a traditional "download the data" model would still have sufficed. However, it was important to provide users the same experience and interfaces they will rely on for later, larger releases, and to get them started in this new paradigm on a small, manageable dataset. DP2, by contrast, marks a new regime, with its size of 59.15 TB outside the range of the traditional download model. The experience users gained working with DP1 through the RSP interfaces can now be applied directly to DP2. The release of DP2 builds on this evolving platform to bring additional upgrades and new functionality, including expanded user query-history and query-status capabilities, context-aware in-application documentation (e.g., contextual help linked directly from Portal table columns), and support for the ADQL spatial `INTERSECTS()` operation.

The RSP remains under active development, with incremental improvements being rolled out as they mature. During the Rubin Early Science Phase, the RSP will continue to operate under a shared-risk model. This section outlines the RSP functionality available at the time of the DP2 release and provides an overview of planned future capabilities.

7.1. *Rubin US Data Access Center*

The Rubin US DAC utilizes a novel hybrid architecture that combines on-premises infrastructure at the USDF at SLAC with flexible, scalable resources in the Google cloud (O'Mullane et al. 2024a). User-facing services are deployed in the cloud to support dynamic scaling in response to user demand and to simplify the provisioning and management of large numbers of science user accounts, whereas the majority of the static data products described in §4 are stored on premises at the USDF, benefiting from cost-effective mass storage and close integration with Rubin's data processing infrastructure, also located at the USDF. For all imaging data except coadds, the Data Butler (§7.2.2) provides the interface between cloud-based users and the on-premises data; for catalog data, a cloud-based TAP client (§7.2.1) submits queries to the on-premises Qserv database cluster (§7.5.1) and retrieves the results. Deep Coadd images are the exception to this hybrid model: because there is one coadd per sky region per band, a relatively small and size-bounded dataset, they are hosted directly in Google Cloud to benefit from cloud-native access.

DP2 is being delivered in two phases because full end-to-end integration of the hybrid RSP for images was not yet complete by the time data processing (§ 5) finished; the hybrid model supporting Qserv (§7.5.1) catalog access over TAP, by contrast, is already in production for EDP2. Rubin therefore adopted a phased release, given that neither coadds nor catalogs depend on the image-side hybrid RSP integration readiness. Early DP2 (EDP2), comprising the deep coadd images and all catalogs, was released on 2026-07-27, while the full DP2 release, adding the visit-level data products once hybrid image integration is demonstrated, is expected in the Oct–Dec 2026 range.

7.2. *Application Programming Interface Aspect*

The Application Programming Interface (API) Aspect provides a comprehensive set of user-facing interfaces for programmatic access to the DP1 data products, through both IVOA-compliant services and the Rubin Data Butler. IVOA services enable standard queries and integration with existing tools, while the Butler facilitates advanced data processing with the LSST Science Pipelines.

Since the DP1 release, additional services have been added as the RSP moves toward full functionality, but some planned services remain unavailable, and certain data products are only accessible via the Butler. This section provides an overview of the available IVOA services and Butler access.

7.2.1. *IVOA Services*

Rubin has adopted a Virtual Observatory (VO)-first design philosophy, prioritizing compliance with IVOA standard interfaces to foster interoperability, standardization, and collaboration. In cases where standardized protocols have yet to be established, additional services

have been introduced to complement these efforts. This approach ensures that the RSP can be seamlessly integrated with community-standard tools such as Tool for OPerations on Catalogues And Tables (TOPCAT) (Taylor 2011) and Aladin (Bonnarel et al. 2000; Boch & Fernique 2014; Baumann et al. 2022), as well as libraries such as PyVO (Graham et al. 2014).

The user-facing APIs are also used internally within the RSP, creating a unified design that ensures consistent and reproducible workflows across all three Aspects. This reduces code duplication, simplifies maintenance, and ensures all users, both internal and external, access data in the same way. For example, an Astronomical Data Query Language (IVOA standard) (ADQL) query on the `Object` catalog via TAP yields identical results whether run from the Portal, Notebook, or an external client.

Since the DP1 release, the RSP has continued to progress towards a full-featured platform, with a number of new services and capabilities added in support of DP2. The following IVOA services are available at the time of the DP2 release.

- **TAP service**: A TAP service (Dowler et al. 2019) enables queries of catalog data via the IVOA-standard ADQL, a dialect of SQL92 with spherical geometry extensions. The main TAP service for DP2 runs on the Rubin-developed Qserv database (§ 7.5), which hosts the core science catalogs described in §4.2, as well as the Visit and Visit Detector catalogs. It also provides image metadata in the IVOA `ObsCore` format via the standard `ivoa.ObsCore` table, making it an "ObsTAP" service (Louys et al. 2017). The TAP service is based on the Canadian Astronomy Data Centre (CADC)'s open-source Java TAP implementation[41], modified for the exact query language accepted by Qserv. Since DP1, ADQL support has been expanded to include the spatial `INTERSECTS()` operation, and temporary-table uploads now provide multi-object query capabilities across both Qserv- and Postgres-backed catalogs. It currently supports a large subset of ADQL, with limitations documented in the Data Release materials (see §8) and exposed via the TAP capabilities endpoint where possible.

  The TAP service provides metadata annotations consistent with the standard, including table and column descriptions, indications of foreign-key relationships between tables, and column metadata such as units and IVOA unified content descriptors (UCDs).

- **Image access services**: Rubin image access services are compliant with IVOA Simple Image Access Protocol, version 2 (SIAv2; Jenness et al. 2024; Dowler et al. 2015) for discovering and accessing astronomical images based on metadata. SIAv2 is a REpresentational State Transfer (REST)-based protocol designed for the discovery and retrieval of image data. It allows, for instance, querying all images in a given band over a defined sky region and time period.

  Users identify an image or observation of interest and query the service. The result set includes metadata about the image, such as the sky position, time, or band, and a data-access URL, which includes an IVOA Identifier uniquely identifying the dataset (Jenness & Dubois-Felsmann 2026), allowing the dataset to be retrieved or a cutout requested via Server-side Operations for Data Access (IVOA standard) (SODA) (SODA; Bonnarel et al. 2017).

- **Image Cutout Service**: The Rubin Cutout Service (Allbery 2023, 2024) is based on the IVOA SODA standard (Bonnarel et al. 2017). Users submit requests specifying sky coordinates and the cutout size as the radius from the coordinates, and the service performs the operation on the full image and returns a result set. As with DP1, the Cutout Service for DP2 remains a single-cutout service, where N cutout requests require N independent synchronous calls; a bulk Cutout Service, (Jenness et al. 2026) is now planned for DR1.

- **HiPS Data Service**: An authenticated HiPS (Fernique et al. 2017) Data Service for seamless pan-and-zoom access to large-scale coadds. It supports fast interactive progressive image exploration at a range of resolutions.

- **Web Distributed Authoring and Versioning (WebDav)**: A WebDav service is provided to enable users to remotely manage, edit, and organize files and directories on the RSP as if they were local files on their own computer. This is especially useful for local development.

Beyond these IVOA services, DP2 introduces expanded user query-history and query-status capabilities, initial context-aware, in-application documentation

[41] https://github.com/opencadc/tap

(e.g., contextual help linked directly from Portal table columns).

Looking ahead to DR1, planned RSP functionality includes bulk cutout services (both for individual targets across a range of epochs and for lists of multiple targets), a PSF retrieval service, an on-demand forced-photometry service, data product recreation services, parallel and batch computing, support for persistent user-created databases and their federation with project-provided databases, support for collaborative work, and full LSSTCam focal-plane visualizations. Access to GPUs and the ability to bring additional user-funded compute resources to the RSP are under consideration for the post-DR1 period.

7.2.2. *Data Butler*

The Rubin Data Butler (Jenness et al. 2022; Lust et al. 2023), is a high-level interface designed to facilitate seamless access to data for both users and software systems. This includes managing storage formats, physical locations, data staging, and database mappings. A Butler repository contains two components.

- The *Data Store*. A physical storage system for datasets, e.g., a Portable Operating System Interface (POSIX) file system or (Amazon) Simple Storage Service (S3) object store.

- The *Registry*. A Structured Query Language (SQL)-compatible database that stores metadata about the datasets in the Data Store.

In the context of the Butler, a dataset refers to a unique data product, such as an image, catalog or map, generated by the observatory or processing pipelines. Datasets belong to one of the various types of data products, described in §4. The Butler ensures that each dataset is uniquely identifiable by a combination of three pieces of information: a data coordinate, a dataset type, and a run collection.

For a DP2 deep coadd on a patch of sky, the dimensions would be the tract, patch and band, specified as a dataId, e.g., `dataId={'tract': 11399, 'patch': 16, 'band': 'g'}` and is associated with the `deep_coadd` DatasetType. In contrast, the `visit_table` and `visit_detector_table` catalogs described in §4 have only the instrument dimension, since each is a single table covering all visits (or all visit-detector combinations) for a given instrument; retrieving either therefore requires only `dataId={'instrument': 'LSSTCam'}`. The `visit_summary` catalog, by contrast, is keyed on band, instrument, `day_obs`, `physical_filter`, and visit, e.g., `dataId={'instrument': 'LSSTCam', 'visit': 2025071100074, 'band': 'g', 'physical_filter': 'g_6', 'day_obs': 20250711}`, and contains one row per detector for that visit.

As for DP1, the EDP2 Butler repository containing the catalogs and coadds described in §4 is hosted in the Google Cloud using an S3-compatible store for datasets and AlloyDB, a PostgreSQL-compatible database for the Registry.

7.2.3. *Remote Programmatic Access*

The Rubin RSP API can be accessed remotely by data rights holders by creating a user security token. This token can then be used as a bearer token for API calls to the RSP services. This capability is especially useful for remote data analysis using tools such as TOPCAT, and for enabling third-party systems, such as community Alert brokers, to access Rubin data. Additionally, it supports remote development, allowing for more flexible workflows and integration with external systems.

7.3. *Portal Aspect*

The Portal Aspect provides an interactive web-based environment for exploratory data discovery, filtering, querying, and visualization of both image and catalog data, without requiring programming expertise. It enables users to access and analyze large datasets via tools for catalog queries, image browsing, time-series inspection, and crossmatching.

The Portal is built on Firefly (Wu et al. 2019), a web application framework developed by the Infrared Processing and Analysis Center. Firefly provides interactive capabilities such as customizable table views, image overlays, multipanel visualizations, and synchronized displays linking catalog and image data.

Designed to support both exploratory data access and detailed scientific investigation, the Portal delivers an intuitive user experience, allowing users to visually analyze data while retaining access to underlying metadata and query controls.

7.4. *Notebook Aspect*

The Notebook Aspect provides an interactive, web-based environment built on Jupyter Notebooks, enabling users to write and execute Python code directly on Rubin and LSST data without downloading it locally. It offers programmatic access to Rubin and LSST data products, allowing users to query and retrieve datasets, manipulate and display images, compute derived properties, plot results, and reprocess data using the LSST Science Pipelines (§5). The environment comes preinstalled with the Pipelines and a broad set of widely used astronomical software tools, supporting immediate and flexible data analysis.

7.5. *Databases*

The user-facing Aspects of the RSP are supported by several backend databases that store catalog data products, image metadata, and other derived datasets. The schema for DP2 and other Rubin databases are available online at https://sdm-schemas.lsst.io/dp2.html.

7.5.1. *Qserv*

The final 10 yr LSST catalog is expected to reach 15 PB and contain measurements for billions of stars and galaxies across trillions of detections. To support efficient storage, querying, and analysis of this dataset, Rubin Observatory developed Qserv (Wang et al. 2011; Mueller et al. 2025), a scalable, parallel, distributed SQL database system. Qserv partitions data over approximately equal-area regions of the celestial sphere, replicates data to ensure resilience and high availability, and uses shared scanning to reduce overall I/O load. It also supports a package of scientific user-defined functions (SciSQL[42]) simplifying complex queries involving spherical geometry, statistics, and photometry. Qserv is built on robust production-quality components, including MariaDB[43] and XRootD[44]. Qserv runs at the USDF and user access to catalog data is via the TAP service (§7.2.1). This enables catalog-based analysis through both the RSP Portal and Notebook Aspects.

While DP1 was small enough that Qserv's full capabilities were not strictly required, DP2, at 59.15 TB, is substantially larger and begins to exercise Qserv at a scale meaningfully closer to its intended operating regime, though still well short of the petabyte-scale catalogs expected for the 10 yr LSST survey.

Qserv is open-source and available on GitHub at https://github.com/lsst/qserv.

7.5.2. *OpSim*

DP2 provides an opsim-format SQLite database of visit-level observational metadata for the visits included in DP2, downloadable from https://s3df.slac.stanford.edu/data/rubin/sim-data/prelsst/dp2_visits.db, with additional documentation available at https://survey-strategy.lsst.io/progress/prelsst/prelsst_dp2.html#visit-database. Like the Visit catalog descibed in §4.5.1, this database contains visit-level entries, but it has no equivalent of the Visit Detector catalog described in §4.5.2 and so provides no per-detector information. Its quantities (e.g., PSF FWHM, 5σ point-source depth) are drawn from the real-time Rapid Analysis QuickLook processing, () performed on the summit rather than from DRP processing. Values may therefore differ somewhat from the corresponding DRP-derived quantities in the Butler `visit_table`. The database also includes some quantities not otherwise available from DRP outputs, such as AOS Zernike coefficients, as well as predicted quantities useful for comparison against simulated survey performance.

8. SUPPORT FOR COMMUNITY SCIENCE

Rubin Observatory's model to support its large, diverse science community to access and analyze DP2 focuses on providing documentation, tutorials, and live virtual learning opportunities alongside an open platform for issue reporting and resolution that enables crowd-sourced solutions. These support resources are all available via the main hub of Rubin's "For Scientists" website[45].

The Data Release documentation for DP2[46] provides an overview of the LSSTCam observations, data quality summary plots, and descriptions of the contents and access methods for all catalog and image datasets. All tutorials (NSF-DOE Vera C. Rubin Observatory 2021) that demonstrate DP2 access via the Rubin Science Platform are delivered in the Data Release documentation, and executable versions of the Jupyter notebook tutorials are also available from a drop-down menu within the Notebook Aspect.

Early DP2 was released on the first day of the annual Rubin Community Workshop, July 27 2026. A dedicated "Get started with DP2" session was added to the Wednesday agenda. Four DP2-onboarding virtual sessions were run in the first week of August, a mix of presentations, tutorials, and time for issue resolution.

The Rubin Community Forum[47] is the open platform for user issue reporting and resolution. Rubin staff monitor all new topics posted in the Support category and ensure that all requests for help get a response and a solution. The open nature of the platform also enables peer-to-peer support that avoids knowledge bottlenecks and hastens the time to resolution.

9. SUMMARY AND FUTURE RELEASES

DP2 provides the first release of science-grade data products from the LSSTCam focal plane. As for DP1, DP2 is part of Rubin's Early Science Program (Guy et al. 2026) and carries a 2 yr proprietary period during

[42] https://smonkewitz.github.io/scisql/
[43] https://www.mariadb.org/
[44] https://xrootd.org/
[45] https://rubinobservatory.org/for-scientists
[46] https://dp2.lsst.io
[47] https://community.lsst.org/

which access is restricted to LSST data rights holders via the cloud-based RSP.

DP2 is delivered in two phases: Early DP2 (EDP2), which provides deep coadded images, all science catalogs, and ancillary data products including throughput curves, survey property maps, and calibration products; and DP2, planned for late 2026, which adds single-epoch visit-level and difference-image data products to complete the release. The coadded images and catalogs delivered in EDP2 will not be reprocessed for DP2; the data are unchanged, and DP2 supersedes and fully encompasses EDP2 by adding the remaining products to the same release. Analyses begun on EDP2 products require no modification.

DP2 represents a substantial advance relative to DP1: the full LSSTCam focal plane replaces the single-raft LSSTComCam, all six *ugrizy* filters are available in place of three, and the sky area is larger by approximately two orders of magnitude. Nevertheless, DP2 reflects the state of an observatory still in the process of optimization: delivered image quality had not yet consistently reached design specification, template coverage is incomplete across the full sky, and several pipeline components remain under active development. Users are encouraged to consult the release documentation at https://dp2.lsst.io and the known issues register maintained there before beginning scientific analyses.

The body of science-grade LSSTCam data acquired between the end of the DP2 observing window on 2026-01-06 and the formal start of the LSST on 2026-06-29 will be processed and made available to LSST data rights holders through the same RSP services. The format and timing of this release are still under discussion; planning will be finalised once DP2 has been delivered.

The first LSST Data Release, DR1, will be derived from a reprocessing of approximately the first year of LSST survey data and is currently expected in mid-2028; its precise definition remains under consideration. DR1 will mark a qualitative advance beyond DP2: observations collected under the community-developed baseline survey strategy (Rubin's Survey Cadence Optimization Committee et al. 2022, 2023, 2025, 2026) will yield substantially broader and more uniform sky coverage across the full WFD footprint of 18,000 $\deg^2$, a balanced six-band filter distribution, and improved delivered image quality as AOS and dome thermal management continue to mature (NSF-DOE Vera C. Rubin Observatory 2026b). Photometric and astrometric calibration will benefit from a full year of all-sky coverage, difference-imaging templates will span a far greater fraction of the footprint and processing will draw on a more mature version of the LSST Science Pipelines refined through DP2 experience and community engagement. Subsequent LSST Data Releases are planned at an approximately yearly cadence, each a complete reprocessing of all survey data accumulated to that point, with progressively deeper coadds and improved calibration. Plans will be updated through revisions to Guy et al. (2026) and on the Rubin community forum at https://community.lsst.org.

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.

This work has been supported by the French National Institute of Nuclear and Particle Physics (IN2P3) through dedicated funding provided by the National Center for Scientific Research (CNRS).

This work has been supported by STFC funding for UK participation in LSST, through grant ST/Y00292X/1.

Facilities: Rubin:Simonyi (LSSTComCam), Rubin:USDAC

Software: Rubin Data Butler (Jenness et al. 2022), LSST Science Pipelines (Rubin Observatory Science Pipelines Developers et al. 2026), LSST Feature Based Scheduler v3.0 (Jones 2026; Yoachim et al. 2024; Naghib et al. 2019), Astropy (Astropy Collaboration et al. 2013, 2018, 2022), PIFF (Jarvis et al. 2021), GBDES (Bernstein 2022), MultiProFit (Taranu 2025), Qserv (Wang et al. 2011; Mueller et al. 2025),

APPENDIX

REFERENCES

Abazajian, K., Adelman-McCarthy, J. K., Agüeros, M. A., et al. 2004, AJ, 128, 502, doi: 10.1086/421365

Acero-Cuellar, T., Bianco, F. B., Sánchez, B. O., Sako, M., & Bellm, E. C. 2026, Performance of Machine-Learned Reliability Scoring for Image Differencing, Data Management Technical Note DMTN-337, NSF-DOE Vera C. Rubin Observatory. https://dmtn-337.lsst.io/

Ahumada, R., Allende Prieto, C., Almeida, A., et al. 2020, ApJS, 249, 3, doi: 10.3847/1538-4365/ab929e

Aihara, H., AlSayyad, Y., Ando, M., et al. 2019, PASJ, 71, 114, doi: 10.1093/pasj/psz103

—. 2022, PASJ, 74, 247, doi: 10.1093/pasj/psab122

Alard, C., & Lupton, R. H. 1998, ApJ, 503, 325, doi: 10.1086/305984

Allbery, R. 2023, IVOA SODA implementation experience, SQuaRE Technical Note SQR-063, NSF-DOE Vera C. Rubin Observatory. https://sqr-063.lsst.io/

—. 2024, Draft IVOA SODA web service specification, SQuaRE Technical Note SQR-093, NSF-DOE Vera C. Rubin Observatory. https://sqr-093.lsst.io/

AlSayyad, Y. 2018, Coaddition Artifact Rejection and CompareWarp, Data Management Technical Note DMTN-080, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2583441

Amouroux, N., Fagrelius, P., Guillemin, T., et al. 2026, arXiv e-prints, arXiv:2606.31822, doi: 10.48550/arXiv.2606.31822

Armstrong, R., Sheldon, E., Huff, E., et al. 2024, arXiv e-prints, arXiv:2407.01771, doi: 10.48550/arXiv.2407.01771

Astier, P., & Regnault, N. 2023, A&A, 670, A118, doi: 10.1051/0004-6361/202245407

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33, doi: 10.1051/0004-6361/201322068

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123, doi: 10.3847/1538-3881/aabc4f

Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, ApJ, 935, 167, doi: 10.3847/1538-4357/ac7c74

Baumann, M., Boch, T., Pineau, F.-X., et al. 2022, in Astronomical Society of the Pacific Conference Series, Vol. 532, Astronomical Data Analysis Software and Systems XXX, ed. J. E. Ruiz, F. Pierfedereci, & P. Teuben, 7

Bechtol, K. 2025, An Interim Report on the On-Sky Commissioning Campaign with LSSTCam, Commissioning Technical Note SITCOMTN-170, NSF-DOE Vera C. Rubin Observatory. https://sitcomtn-170.lsst.io/

Berk, A., Anderson, G. P., Bernstein, L. S., et al. 1999, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 3756, Optical Spectroscopic Techniques and Instrumentation for Atmospheric and Space Research III, ed. A. M. Larar, 348–353, doi: 10.1117/12.366388

Bernstein, G., Gruendl, R., & Kent, S. 2016, Limiting Magnitude, and Image Quality in DES Year 1, Tech. rep., Fermi National Accelerator Laboratory (FNAL), Batavia, IL (United States), doi: 10.2172/1250877

Bernstein, G. M. 2022, gbdes: DECam instrumental signature fitting and processing programs, Astrophysics Source Code Library, record ascl:2210.011. http://ascl.net/2210.011

Bernstein, G. M., Armstrong, R., Plazas, A. A., et al. 2017, PASP, 129, 074503, doi: 10.1088/1538-3873/aa6c55

Bertin, E. 2011, in Astronomical Society of the Pacific Conference Series, Vol. 442, Astronomical Data Analysis Software and Systems XX, ed. I. N. Evans, A. Accomazzi, D. J. Mink, & A. H. Rots, 435

Bianco, F. B., Ivezić, Ž., Jones, R. L., et al. 2022, ApJS, 258, 1, doi: 10.3847/1538-4365/ac3e72

Blum, R., & the Rubin Operations Team. 2020, Vera C. Rubin Observatory Data Policy, Data Management Operations Controlled Document RDO-013, NSF-DOE Vera C. Rubin Observatory. https://ls.st/RDO-013

Boch, T., & Fernique, P. 2014, in Astronomical Society of the Pacific Conference Series, Vol. 485, Astronomical Data Analysis Software and Systems XXIII, ed. N. Manset & P. Forshay, 277

Bonnarel, F., Dowler, P., Demleitner, M., Tody, D., & Dempsey, J. 2017, IVOA Server-side Operations for Data Access Version 1.0, IVOA Recommendation 17 May 2017, doi: 10.5479/ADS/bib/2017ivoa.spec.0517B

Bonnarel, F., Fernique, P., Bienaymé, O., et al. 2000, A&AS, 143, 33, doi: 10.1051/aas:2000331

Bosch, J., Armstrong, R., Bickerton, S., et al. 2018, PASJ, 70, S5, doi: 10.1093/pasj/psx080

Bosch, J. F. 2026, Challenges serializing Rubin Observatory data products, Data Management Technical Note DMTN-339, NSF-DOE Vera C. Rubin Observatory. https://dmtn-339.lsst.io/

Broughton, A., Ferté, A., Waters, C. Z., et al. 2026, Image Calibration and Instrument Signal Removal for the First Year of the LSST, Technical Note RTN-117, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/3488007

Broughton, A., Utsumi, Y., Plazas Malagón, A. A., et al. 2024, PASP, 136, 045003, doi: 10.1088/1538-3873/ad3aa2

Burke, D. L., Rykoff, E. S., Allam, S., et al. 2018, AJ, 155, 41, doi: 10.3847/1538-3881/aa9f22

Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv e-prints, arXiv:1612.05560, doi: 10.48550/arXiv.1612.05560

Claver, C. F., & The LSST Systems Engineering Integrated Project Team. 2017, LSST System Requirements (LSR), Systems Engineering Controlled Document LSE-29, NSF-DOE Vera C. Rubin Observatory. https://ls.st/LSE-29

Corbett, H., Carney, J., Marshall, W., et al. 2024, in Software and Cyberinfrastructure for Astronomy VIII, ed. J. Ibsen & G. Chiozzi, Vol. 13101, International Society for Optics and Photonics (SPIE), 1310137, doi: 10.1117/12.3018987

Dowler, P., Bonnarel, F., & Tody, D. 2015, IVOA Simple Image Access Version 2.0, IVOA Recommendation 23 December 2015, doi: 10.5479/ADS/bib/2015ivoa.spec.1223D

Dowler, P., Rixon, G., Tody, D., & Demleitner, M. 2019, Table Access Protocol Version 1.1, IVOA Recommendation 27 September 2019, doi: 10.5479/ADS/bib/2019ivoa.spec.0927D

Drlica-Wagner, A., Taranto, A., Rodeghiero, G., et al. 2026, Investigation and Mitigation of a Prominent Off-Axis Stray Light Path in Rubin Observatory Commissioning. https://arxiv.org/abs/2606.31939

Ellis, K. S. 2009, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7427, Optical Modeling and Performance Predictions IV, ed. M. A. Kahan, 742708, doi: 10.1117/12.830599

Euclid Collaboration, Aussel, H., Tereno, I., et al. 2026a, A&A, 711, A1, doi: 10.1051/0004-6361/202554610

Euclid Collaboration, Romelli, E., Kümmel, M., et al. 2026b, A&A, 711, A4, doi: 10.1051/0004-6361/202554586

Euclid Collaboration, Matamoro Zatarain, T., Fotopoulou, S., et al. 2026c, A&A, 711, A20, doi: 10.1051/0004-6361/202554619

Fagrelius, P., & Rykoff, E. S. 2025, Rubin Observatory Baseline Calibration Plan, Commissioning Technical Note SITCOMTN-086, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2583850

Ferguson, P. S., Rykoff, E. S., Carlin, J. L., Saunders, C., & Parejko, J. K. 2025, The Monster: A reference catalog with synthetic ugrizy-band fluxes for the Vera C. Rubin observatory, Data Management Technical Note DMTN-277, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2583688

Fernique, P., Allen, M., Boch, T., et al. 2017, HiPS - Hierarchical Progressive Survey Version 1.0, IVOA Recommendation 19 May 2017, doi: 10.5479/ADS/bib/2017ivoa.spec.0519F

Fortino, W. F., Bernstein, G. M., Bernardinelli, P. H., et al. 2021, AJ, 162, 106, doi: 10.3847/1538-3881/ac0722

Gaia Collaboration, Montegriffo, P., Bellazzini, M., et al. 2023, A&A, 674, A33, doi: 10.1051/0004-6361/202243709

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759, doi: 10.1086/427976

Graham, M., Plante, R., Tody, D., & Fitzpatrick, M. 2014, PyVO: Python access to the Virtual Observatory, Astrophysics Source Code Library, record ascl:1402.004

Guy, L. P., AlSayyad, Y., Bechtol, K., et al. 2026, Rubin Observatory Plans for an Early Science Program, Technical Note RTN-011, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2584021

Holman, M. J., Bernardinelli, P. H., Schwamb, M. E., et al. 2025, AJ, 170, 97, doi: 10.3847/1538-3881/ade435

Howard, J., Reil, K., Claver, C., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10700, Ground-based and Airborne Telescopes VII, ed. H. K. Marshall & J. Spyromilio, 107003D, doi: 10.1117/12.2312684

Ivezić, Ž., & The LSST Science Collaboration. 2018, LSST Science Requirements Document, Project Controlled Document LPM-17, NSF-DOE Vera C. Rubin Observatory. https://ls.st/LPM-17

Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019a, ApJ, 873, 111, doi: 10.3847/1538-4357/ab042c

—. 2019b, ApJ, 873, 111, doi: 10.3847/1538-4357/ab042c

Jarvis, M., Bernstein, G. M., Amon, A., et al. 2021, MNRAS, 501, 1282, doi: 10.1093/mnras/staa3679

Jenness, T., & Dubois-Felsmann, G. P. 2026, IVOA Identifier Usage at the Rubin Observatory, Data Management Technical Note DMTN-302, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2583848

Jenness, T., Dubois-Felsmann, G. P., & Bosch, J. F. 2026, Bulk Cutout Service Implementation Options, Data Management Technical Note DMTN-326, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/3399549

Jenness, T., Voutsinas, S., Dubois-Felsmann, G. P., & Salnikov, A. 2024, arXiv e-prints, arXiv:2501.00544, doi: 10.48550/arXiv.2501.00544

Jenness, T., Bosch, J. F., Salnikov, A., et al. 2022, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12189, Software and Cyberinfrastructure for Astronomy VII, 1218911, doi: 10.1117/12.2629569

Jones, R. L. 2026, Scheduling the LSST: Commissioning the Rubin Feature Based Scheduler, Simulations Team Technical Note SMTN-020, NSF-DOE Vera C. Rubin Observatory. https://smtn-020.lsst.io/

Jones, R. L., Yoachim, P., Neilsen, Eric H., J., et al. 2026, in Proceedings of SPIE, Vol. 14151, Observatory Operations: Strategies, Processes, and Systems XI, 141510X, doi: 10.1117/12.3104048

Jurić, M., Ciardi, D., Dubois-Felsmann, G., & Guy, L. 2019, LSST Science Platform Vision Document, Systems Engineering Controlled Document LSE-319, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2587242

Kannawadi, A. 2025, Consistent galaxy colors with Gaussian-Aperture and PSF photometry, Data Management Technical Note DMTN-190, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2583849

Karpov, S., & Peloton, J. 2022, arXiv e-prints, arXiv:2202.05719, doi: 10.48550/arXiv.2202.05719

Kron, R. G. 1980, ApJS, 43, 305, doi: 10.1086/190669

Kuijken, K. 2008, A&A, 482, 1053, doi: 10.1051/0004-6361:20066601

Laliotis, K., Léget, P.-F., Roodman, A., et al. 2026, Impact of CCD Non-Flatness In Rubin LSSTCam PSF Measurements and Modeling, Technical Note RTN-108, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/3399537

Léget, P.-F. 2026, Point Spread Function Characterization of the Vera C. Rubin Observatory using the Data Preview 2 dataset., Technical Note RTN-126, NSF-DOE Vera C. Rubin Observatory. https://rtn-126.lsst.io/

Léget, P. F., Astier, P., Regnault, N., et al. 2021, A&A, 650, A81, doi: 10.1051/0004-6361/202140463

Louys, M., Tody, D., Dowler, P., et al. 2017, Observation Data Model Core Components, its Implementation in the Table Access Protocol Version 1.1, IVOA Recommendation 09 May 2017, doi: 10.5479/ADS/bib/2017ivoa.spec.0509L

LSST Camera Team, Antilogus, P., Astier, P., et al. 2025, LSST Camera Electro-Optical Test (Run 7) Results, Commissioning Technical Note SITCOMTN-148, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2583999

LSST Dark Energy Science Collaboration (LSST DESC), Abolfathi, B., Alonso, D., et al. 2021, ApJS, 253, 31, doi: 10.3847/1538-4365/abd62c

Lupton, R., Blanton, M. R., Fekete, G., et al. 2004, PASP, 116, 133, doi: 10.1086/382245

Lust, N. B., Jenness, T., Bosch, J. F., et al. 2023, arXiv e-prints, arXiv:2303.03313, doi: 10.48550/arXiv.2303.03313

Megias Homar, G., Kahn, S. M., Meyers, J. M., Crenshaw, J. F., & Thomas, S. J. 2024, ApJ, 974, 108, doi: 10.3847/1538-4357/ad6cdc

Melchior, P., Moolekamp, F., Jerdee, M., et al. 2018, Astronomy and Computing, 24, 129, doi: 10.1016/j.ascom.2018.07.001

Merritt, S. R., Fedorets, G., Schwamb, M. E., et al. 2025, AJ, 170, 100, doi: 10.3847/1538-3881/add3ec

Mueller, F., Gaponenko, I., Gates, J., et al. 2025, in Astronomical Society of the Pacific Conference Series, Vol. 538, Astronomical Data Analysis Software and Systems XXXII, ed. S. Gaudet, D. Bohlender, S. Gwyn, A. Hincks, & P. Teuben, 114, doi: 10.26624/XCPI7375

Muinonen, K., Belskaya, I. N., Cellino, A., et al. 2010, Icarus, 209, 542, doi: 10.1016/j.icarus.2010.04.003

Naghib, E., Yoachim, P., Vanderbei, R. J., Connolly, A. J., & Jones, R. L. 2019, AJ, 157, 151, doi: 10.3847/1538-3881/aafece

NSF-DOE Vera C. Rubin Observatory. 2021, Rubin Observatory LSST Tutorials [Computer Software], NSF-DOE Vera C. Rubin Observatory, doi: 10.11578/rubin/dc.20250909.20

—. 2025, Legacy Survey of Space and Time Data Preview 1 [Data set], NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/RUBIN/2570308

—. 2026a, Rubin Observatory Early Operations On-sky Performance, Technical Note RTN-122, NSF-DOE Vera C. Rubin Observatory. https://rtn-122.lsst.io/

—. 2026b, Image Quality Improvement Plan, Summit Operations Technical Note SOTN-004, NSF-DOE Vera C. Rubin Observatory. https://sotn-004.lsst.io/

Oke, J. B., & Gunn, J. E. 1983, ApJ, 266, 713, doi: 10.1086/160817

O'Mullane, W., AlSayyad, Y., Chiang, J., et al. 2024a, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 13101, Software and Cyberinfrastructure for Astronomy VIII, ed. J. Ibsen & G. Chiozzi, 131012B, doi: 10.1117/12.3018005

O'Mullane, W., Economou, F., Huang, F., et al. 2024b, in Astronomical Society of the Pacific Conference Series, Vol. 535, Astromical Data Analysis Software and Systems XXXI, ed. B. V. Hugo, R. Van Rooyen, & O. M. Smirnov, 227, doi: 10.48550/arXiv.2111.15030

Onken, C. A., Wolf, C., Bessell, M. S., et al. 2019, PASA, 36, e033, doi: 10.1017/pasa.2019.27

Pai, A., Léget, P.-F., & Chang, C. 2026, Impact of Including Wavelength-dependent PSF Modeling in Rubin Commissioning Data, Commissioning Technical Note SITCOMTN-174, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/3398709

Petrosian, V. 1976, ApJL, 210, L53, doi: 10.1086/18230110.1086/182253

Plazas Malagón, A. A., Digel, S. W., Roodman, A., et al. 2026, LSSTCam and LSSTComCam Focal Plane Layouts, Camera Technical Note CTN-001, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2584019

Plazas Malagón, A. A., Waters, C., Broughton, A., et al. 2025, Journal of Astronomical Telescopes, Instruments, and Systems, 11, 011209, doi: 10.1117/1.JATIS.11.1.011209

Pollek, H. M. M., Rodeghiero, G., Andrew, J., et al. 2026, Mechanical Studies of an Additional Light Baffle for the LSST Camera. https://arxiv.org/abs/2606.31898

Reiss, D. J., & Lupton, R. H. 2016, Implementation of Image Difference Decorrelation, Data Management Technical Note DMTN-021, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2586490

Rodeghiero, G., Drlica-Wagner, A., Taranto, A., et al. 2026, An overview of stray light findings and interpretation during on-sky commissioning of LSSTCam. https://arxiv.org/abs/2606.31945

Rubin, V. C., & Ford, Jr., W. K. 1970, ApJ, 159, 379, doi: 10.1086/150317

Rubin, V. C., Ford, Jr., W. K., & Thonnard, N. 1980, ApJ, 238, 471, doi: 10.1086/158003

Rubin Observatory Science Pipelines Developers, AlSayyad, Y., Armstrong, R., et al. 2026, The LSST Science Pipelines Software: Optical Survey Pipeline Reduction and Analysis Environment, Project Science Technical Note PSTN-019, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2570545

Rubin's Survey Cadence Optimization Committee, Bauer, F. E., Brough, S., et al. 2022, Survey Cadence Optimization Committee's Phase 1 Recommendation, Project Science Technical Note PSTN-053, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2584276

Rubin's Survey Cadence Optimization Committee, Bauer, F. E., Bianco, F. B., et al. 2023, Survey Cadence Optimization Committee's Phase 2 Recommendations, Project Science Technical Note PSTN-055, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2585249

Rubin's Survey Cadence Optimization Committee, Bianco, F. B., Jones, R. L., et al. 2025, Survey Cadence Optimization Committee's Phase 3 Recommendations, Project Science Technical Note PSTN-056, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2585402

—. 2026, Survey Cadence Optimization Committee's Survey Start Recommendations, Project Science Technical Note PSTN-057, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/3395070

Rykoff, E. S., Tucker, D. L., Burke, D. L., et al. 2023, arXiv e-prints, arXiv:2305.01695, doi: 10.48550/arXiv.2305.01695

Saunders, C. 2024, Astrometric Calibration in the LSST Pipeline, Data Management Technical Note DMTN-266, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2583846

—. 2025, Early astrometric residuals characterization of LSSTCam: Atmospheric turbulences and static sensor effect., Data Management Technical Note DMTN-324, NSF-DOE Vera C. Rubin Observatory. https://dmtn-324.lsst.io/

Schutt, T., Jarvis, M., Roodman, A., et al. 2025, The Open Journal of Astrophysics, 8, 26, doi: 10.33232/001c.132299

Sérsic, J. L. 1963, Boletin de la Asociacion Argentina de Astronomia La Plata Argentina, 6, 41

Sersic, J. L. 1968, Atlas de Galaxias Australes (Cordoba, Argentina: Observatorio Astronomico)

Shanks, T., Metcalfe, N., Chehade, B., et al. 2015, MNRAS, 451, 4238, doi: 10.1093/mnras/stv1130

Sheldon, E. S., Becker, M. R., MacCrann, N., & Jarvis, M. 2020, ApJ, 902, 138, doi: 10.3847/1538-4357/abb595

Shuntov, M., Akins, H. B., Paquereau, L., et al. 2025, A&A, 704, A339, doi: 10.1051/0004-6361/202555799

SLAC National Accelerator Laboratory, & NSF-DOE Vera C. Rubin Observatory. 2024, LSST Commissioning Camera, SLAC National Accelerator Laboratory (SLAC), Menlo Park, CA (United States), doi: 10.71929/RUBIN/2561361

—. 2025, The LSST Camera (LSSTCam), SLAC National Accelerator Laboratory (SLAC), Menlo Park, CA (United States), doi: 10.71929/rubin/2571927

Stalder, B., Reil, K., Claver, C., et al. 2020, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 11447, Ground-based and Airborne Instrumentation for Astronomy VIII, ed. C. J. Evans, J. J. Bryant, & K. Motohara, 114470L, doi: 10.1117/12.2561132

Stalder, B., Reil, K., Aguilar, C., et al. 2022, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12184, Ground-based and Airborne Instrumentation for Astronomy IX, ed. C. J. Evans, J. J. Bryant, & K. Motohara, 121840J, doi: 10.1117/12.2630184

Stalder, B., Munoz, F., Aguilar, C., et al. 2024, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 13094, Ground-based and Airborne Telescopes X, ed. H. K. Marshall, J. Spyromilio, & T. Usuda, 1309409, doi: 10.1117/12.3019266

Taranto, A., Rodeghiero, G., Rosignoli, L., et al. 2026, Modeling of the diffuse background produced by the Vera C. Rubin Observatory M2 baffle scattered light. https://arxiv.org/abs/2606.31793

Taranu, D. S. 2025, The MultiProFit astronomical source modelling code, Data Management Technical Note DMTN-312, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2584108

Taylor, M. 2011, TOPCAT: Tool for OPerations on Catalogues And Tables, Astrophysics Source Code Library, record ascl:1101.010

Thomas, S., Connolly, A., Crenshaw, J. F., et al. 2023, in Adaptive Optics for Extremely Large Telescopes (AO4ELT7), 67, doi: 10.13009/AO4ELT7-2023-069

Vera C. Rubin Observatory Team. 2026, The Vera C. Rubin Observatory Data Preview 1, Technical Note RTN-095, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2570536

Vera C. Rubin Observatory Team, Acero-Cuellar, T., Acosta, E., et al. 2026, AJ, 171, 360, doi: 10.3847/1538-3881/ae521f

Wang, D. L., Monkewitz, S. M., Lim, K.-T., & Becla, J. 2011, in State of the Practice Reports, SC '11 (New York, NY, USA: ACM), 12:1–12:11, doi: 10.1145/2063348.2063364

Wu, X., Roby, W., Goldian, T., et al. 2019, in Astronomical Society of the Pacific Conference Series, Vol. 521, Astronomical Data Analysis Software and Systems XXVI, ed. M. Molinaro, K. Shortridge, & F. Pasian, 32

Xin, B., Claver, C., Liang, M., et al. 2015, ApOpt, 54, 9045, doi: 10.1364/AO.54.009045

Yoachim, P., Jones, L., Eric H. Neilsen, J., & Becker, M. R. 2024, lsst/rubin_scheduler: v3.0.0, v3.0.0, Zenodo, doi: 10.5281/zenodo.13985198

Zuntz, J., Sheldon, E., Samuroff, S., et al. 2018, MNRAS, 481, 1149, doi: 10.1093/mnras/sty2219

Glossary

2D: Two-dimensional.

ADQL: Astronomical Data Query Language (IVOA standard).

ADU: Analogue-to-Digital Unit.

airmass: The pathlength of light from an astrophysical source through the Earth's atmosphere. It is given approximately by sec z, where z is the angular distance from the zenith (the point directly overhead, where airmass = 1.0) to the source.

algorithm: A computational implementation of a calculation or some method of processing.

AOS: Active Optics System.

API: Application Programming Interface.

astrometry: In astronomy, the sub-discipline of astrometry concerns precision measurement of positions (at a reference epoch), and real and apparent motions of astrophysical objects. Real motion means 3-D motions of the object with respect to an inertial reference frame; apparent motions are an artifact of the motion of the Earth. Astrometry per se is sometimes confused with the act of determining a World Coordinate System (WCS), which is a functional characterization of the mapping from pixels in an image or spectrum to world coordinate such as (RA, Dec) or wavelength.

ATLAS: Asteroid Terrestrial-impact Last Alert System.

au: astronomical unit.

AURA: Association of Universities for Research in Astronomy.

AuxTel: LSST's 1.2-meter Auxiliary Telescope will measure atmospheric transmission and will be used to calibrate LSST images.

background: In an image, the background consists of contributions from the sky (e.g., clouds or scattered moonlight), and from the telescope and camera optics, which must be distinguished from the astrophysical background. The sky and instrumental backgrounds are characterized and removed by the LSST processing software using a low-order spatial function whose coefficients are recorded in the image metadata.

BNL: Brookhaven National Laboratory.

Butler: A middleware component for persisting and retrieving image datasets (raw or processed), calibration reference data, and catalogs.

CADC: Canadian Astronomy Data Centre.

calibration: The process of translating signals produced by a measuring instrument such as a telescope and camera into physical units such as flux, which are used for scientific analysis. Calibration removes most of the contributions to the signal from environmental and instrumental factors, such that only the astronomical component remains.

camera: An imaging device mounted at a telescope focal plane, composed of optics, a shutter, a set of filters, and one or more sensors arranged in a focal plane array.

CBP: Collimated Beam Projector.

CCD: Charge-Coupled Device.

cloud: A visible mass of condensed water vapor floating in the atmosphere, typically high above the ground or in interstellar space acting as the birthplace for stars. Also a way of computing (on other peoples computers leveraging their services and availability).

configuration: A task-specific set of configuration parameters, also called a 'config'. The config is read-only; once a task is constructed, the same configuration will be used to process all data. This makes the data processing more predictable: it does not depend on the order in which items of data are processed. This is distinct from arguments or options, which are allowed to vary from one task invocation to the next.

COSMOS: Cosmic Evolution Survey.

CTI: Charge Transfer Inefficiency.

DAC: Data Access Center.

DAQ: Data Acquisition System.

Data Management System: The computing infrastructure, middleware, and applications that process, store, and enable information extraction from the LSST dataset; the DMS will process peta-scale data volume, convert raw images into a faithful representation of the universe, and archive the results in a useful form. The infrastructure layer consists of the computing, storage, networking hardware, and system software. The middleware layer handles distributed processing, data access, user interface, and system operations services. The applications layer includes the data pipelines and the science data archives' products and services.

Data Release Processing: Deprecated term; see Data Release Production.

DCR: Differential Chromatic Refraction.

DDF: Deep Drilling Field.

deblend: Deblending is the act of inferring the intensity profiles of two or more overlapping sources from a single footprint within an image. Source footprints may overlap in crowded fields, or where the astrophysical phenomena intrinsically overlap (e.g., a supernova embedded in an external galaxy), or by spatial co-incidence (e.g., an asteroid passing in front of a star). Deblending may make use of a priori information from images (e.g., deep CoAdds or visit images obtained in good seeing), from catalogs, or from models. A 'deblend' is commonly referred to in terms of 'parent' (total) and 'child' (component) objects.

DECam: Dark Energy Camera.

deg: degree; unit of angle.

DES: Dark Energy Survey.

DIA: Difference Image Analysis.

DMS: Data Management Subsystem.

DOE: Department of Energy.

DP0: Data Preview 0.

DP1: Data Preview 1.

DP2: Data Preview 2.

DR1: Data Release 1.

DR3: Data Release 3.

DRP: Data Release Processing.

E2V: Teledyne-e2v semiconductor company.

ECDFS: Extended Chandra Deep Field-South Survey.

EDFS: Euclid Deep Field South.

EDP2: Early Data Preview 2.

EPO: Education and Public Outreach.

epoch: Sky coordinate reference frame, e.g., J2000. Alternatively refers to a single observation (usually photometric, can be multi-band) of a variable source.

exponential profile: The radial distribution of flux of an astronomical source that is characterized: I(r)=I0exp(.68(r/re)) The normalization 1.68 is chosen so that the model radius is a half-light radius. An 2-dimensional elliptical version of this profile is fit to every detected source.

FBS: Feature-Based Scheduler.

FGCM: Forward Global Calibration Method.

Firefly: A framework of software components written by IPAC for building web-based user interfaces to astronomical archives, through which data may be searched and retrieved, and viewed as FITS images, catalogs, and/or plots. Firefly tools will be integrated into the Science Platform.

flux: Shorthand for radiative flux, it is a measure of the transport of radiant energy per unit area per unit time. In astronomy this is usually expressed in cgs units: erg/cm2/s.

footprint: See 'source footprint', 'instrumental footprint', or 'survey footprint', 'Footprint' is a Python class representing a source footprint.

forced photometry: A measurement of the photometric properties of a source, or expected source, with one or more parameters held fixed. Most often this means fixing the location of the center of the brightness profile (which may be known or predicted in advance), and measuring other properties such as total brightness, shape, and orientation. Forced photometry will be done for all Objects in the Data Release Production.

FWHM: Full Width at Half-Maximum.

GAaP: Gaussian Aperture and PSF.

Gaia: a space observatory of the European Space Agency, launched in 2013 and expected to operate until 2025. The spacecraft is designed for astrometry: measuring the positions, distances and motions of stars with unprecedented precision.

HEALPix: Hierarchical Equal-Area iso-Latitude Pixelisation.

HiPS: Hierarchical Progressive Survey (IVOA standard).

HSM: Shape measurement algorithm from Hirata & Seljak (2003) and Mandelbaum et al. (2005).

IAU: International Astronomical Union.

ISR: Instrument Signal Removal.

ITL: Imaging Technology Laboratory (UA).

IVOA: International Virtual Observatory Alliance.

L1: Lens 1.

L2: Lens 2.

L3: Lens 3.

LED: Light-Emitting Diode.

LSST: Legacy Survey of Space and Time.

LSSTCam: LSST Science Camera.

LSSTComCam: Rubin Commissioning Camera.

LWS: Light-Wind Screen.

M1: Primary Mirror.

M1M3: Single piece of glass for Primary Mirror/Tertiary Mirror.

M2: Secondary Mirror.

metadata: General term for data about data, e.g., attributes of astronomical objects (e.g. images, sources, astroObjects, etc.) that are characteristics of the objects themselves, and facilitate the organization, preservation, and query of data sets (e.g., a FITS header contains metadata).

MJD: Modified Julian Date.

MPC: Minor Planet Center.

MPS: NSF Mathematical and Physical Sciences directorate.

MultiProFit: The MultiProFit astronomical source modelling code, including the multiprofit Python package and its dependencies, and the meas_extensions_multiprofit package for pipeline tasks.

NIST: National Institute of Standards and Technology (USA).

NSF: National Science Foundation.

Object: In LSST nomenclature this refers to an astronomical object, such as a star, galaxy, or other physical entity. E.g., comets, asteroids are also Objects but typically called a Moving Object or a Solar System Object (SSObject). One of the DRP data products is a table of Objects detected by LSST which can be static, or change brightness or position with time.

Pan-STARRS: Panoramic Survey Telescope and Rapid Response System.

passband: The window of wavelength or the energy range admitted by an optical system; specifically the transmission as a function of wavelength or energy. Typically the passband is limited by a filter. The width of the passband may be characterized in a variety of ways, including the width of the half-power points of the transmission curve, or by the equivalent width of a filter with 100% transmission within the passband, and zero elsewhere.

patch: An quadrilateral sub-region of a sky tract, with a size in pixels chosen to fit easily into memory on desktop computers.

pipeline: A configured sequence of software tasks (Stages) to process data and generate data products. Example: Association Pipeline.

POSIX: Portable Operating System Interface.

PSF: Point Spread Function.

PTC: Photon Transfer Curve.

PVI: Processed Visit Image.

QE: quantum efficiency.

Qserv: LSST's distributed parallel database. This database system is used for collecting, storing, and serving LSST Data Release Catalogs and Project metadata, and is part of the Software Stack.

REB: Readout Electronics Board.

REST: REpresentational State Transfer.

RSP: Rubin Science Platform.

S3: (Amazon) Simple Storage Service.

scarlet lite: The lightweight version of the scarlet algorithm, optimized for multiband deblending on pixel-aligned grids. Includes the scarlet_lite and meas_extensions_scarlet Python packages.

schema: The definition of the metadata and linkages between datasets and metadata entities in a collection of data or archive.

SDSS: Sloan Digital Sky Survey.

shape: In reference to a Source or Object, the shape is a functional characterization of its spatial intensity distribution, and the integral of the shape is the flux. Shape characterizations are a data product in the DIASource, DIAObject, Source, and Object catalogs.

ShearObject: An object with measurements using the metadetection algorithm, or the catalog thereof.

Simonyi Survey Telescope: The telescope at the Rubin Observatory that will perform the LSST (this refers to all physical components: the mirror, the mount assembly, etc.).

SLAC: SLAC National Accelerator Laboratory.

SLR: Stellar Locus Regression.

SNR: Signal to Noise Ratio.

SODA: Server-side Operations for Data Access (IVOA standard).

software: The programs and other operating information used by a computer.

Solar System Object: A solar system object is an astrophysical object that is identified as part of the Solar System: planets and their satellites, asteroids, comets, etc. This class of object had historically been referred to within the LSST Project as Moving Objects.

Source: A single detection of an astrophysical object in an image, the characteristics for which are stored in the Source Catalog of the DRP database. The association of Sources that are non-moving lead to Objects; the association of moving Sources leads to Solar System Objects. (Note that in non-LSST usage 'source' is often used for what LSST calls an Object.).

SQL: Structured Query Language.

SV: Science Validation.

TAI: International Atomic Time.

TAP: Table Access Protocol (IVOA standard).

TBD: To Be Defined (Determined).

TOPCAT: Tool for OPerations on Catalogues And Tables.

tract: A portion of sky, a spherical convex polygon, within the LSST all-sky tessellation (sky map). Each tract is subdivided into sky patches.

transient: A transient source is one that has been detected on a difference image, but has not been associated with either an astronomical object or a solar system body.

USDF: United States Data Facility.

Visit: A sequence of one or more consecutive exposures at a given position, orientation, and filter within the LSST cadence. See Standard Visit, Alternate Standard Visit, and Non-Standard Visit.

VO: Virtual Observatory.

VST: VLT Survey Telescope.

WCS: World Coordinate System.

WebDav: Web Distributed Authoring and Versioning.

WFD: Wide-Fast-Deep.

XP: B or R Photometry (Gaia).