Draft version June 29, 2025
Typeset using LaTeX twocolumn style in AASTeX7

https://doi.org/10.71929/rubin/2570536

The Vera C. Rubin Observatory Data Preview 1

(The Vera C. Rubin Observatory)

(Dated: June 29, 2025)

ABSTRACT

We present Data Preview 1 (DP1), the first data from the National Science Foundation ()-Department of Energy () Vera C. Rubin Observatory, comprising raw and calibrated single-epoch images, coadds, difference images, detection catalogs, and derived data products. DP1 is based on 1792 science-grade optical/near-infrared exposures acquired over 48 distinct nights by the Rubin Commissioning Camera, LSSTComCam, on the Simonyi Survey Telescope at the Summit Facility on Cerro Pachón, Chile during the first on-sky commissioning campaign in late 2024. DP1 covers ∼15 sq. deg. over seven roughly equally-sized non-contiguous fields, each independently observed in six broad photometric bands, *ugrizy*, spanning a range of stellar densities and latitudes and overlapping with external reference datasets. The median image quality across all bands, measured by the Full Width at Half-Maximum (FWHM) of the point-spread function, is approximately 1.13 arcseconds, with the sharpest images reaching about 0.65 arcseconds. DP1 contains approximately 2.3 million distinct astrophysical objects, of which 1.6 million are extended in at least one band, and 431 solar system objects, of which 93 are new discoveries. DP1 is approximately 3.5 TB in size and available to Rubin data rights holders via the Rubin Science Platform, a cloud-based environment for the analysis of petascale astronomical data. While small compared to future Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope) (LSST) releases, its high quality and diversity of data support a broad range of early science investigations across all four LSST themes ahead of full operations in late 2025.

Keywords: Rubin Observatory – LSST

1. INTRODUCTION

The NSF–DOE Vera C. Rubin Observatory is a ground-based, wide-field optical/near-infrared facility located on Cerro Pachón in northern Chile. Named in honor of Vera C. Rubin, a pioneering astronomer whose groundbreaking work in the 20th century provided the first convincing evidence for the existence of dark matter (Rubin & Ford 1970; Rubin et al. 1980), the observatory's prime mission is to carry out the LSST (Ivezić et al. 2019a). This 10-year survey is designed to obtain rapid-cadence, multi-band imaging of the entire visible southern sky approximately every 3–4 nights, mapping it to a depth of $\sim$ 27.5 magnitude in the r-band with ∼0.7 arcsecond seeing, with a total of ∼800 visits per pointing.

The Rubin Observatory system consists of four main components: the Simonyi Survey Telescope, featuring an 8.4 m diameter (6.5 m effective aperture) primary mirror that delivers a wide field of view; a 3.2-gigapixel Camera, capable of imaging 9.6 square degrees per exposure with seeing-limited quality in six broadband filters, *ugrizy* (320–1050 nm); an automated Data Management System that processes and archives tens of terabytes of data per night, generating science-ready data products within minutes for a global community of scientists; and an Education and Public Outreach () program that provides real-time data access, interactive tools, and educational content to engage the public. The integrated system's étendue[1] of 319 $\mathrm{m}^2\,\mathrm{deg}^2$, is over an order of magnitude larger than that of any existing facility, enabling a fast, large-scale survey with exceptional depth in a fraction of the time compared to other observatories.

[1] The product of the primary mirror area and the angular area of its field of view for a given set of observing conditions.

The observatory's design is driven by four key science themes: probing dark energy and dark matter; taking an inventory of the solar system; exploring the transient optical sky; and mapping the Milky Way (Ivezić et al. 2019a). These themes inform the optimization of a range of system parameters, including image quality, photometric and astrometric accuracy, the depth of a single visit and the co-added survey depth, the filter complement, the total number of visits per pointing as well as the distribution of visits on the sky, and total sky coverage. Additionally, they inform the design of the data processing and access systems. By optimizing the system parameters to support a wide range of scientific goals, we maximize the observatory's scientific output across all areas, transforming Rubin into a powerful discovery machine capable of addressing a broad range of astrophysical questions.

Over the lifetime of the LSST, Rubin Observatory will issue several Data Releases, each representing a full reprocessing of all LSST data collected to date. Prior to the start of the LSST survey, commissioning activities will generate a significant volume of science-grade data. To make this early data available to the community, the Rubin Early Science Program, (Guy et al. 2025), was established. One key component of this program is a series of Data Previews; early versions of the LSST Data Releases. These previews include preliminary data products derived from both simulated and commissioning data, which, together with early versions of the data access services, are intended to support high-impact early science, facilitate community readiness, and inform the development of Rubin's operational capabilities ahead of the start of full survey operations. All data and services provided through the Rubin Early Science Program are offered on a shared-risk basis[2].

This paper describes Rubin's second of three planned Data Previews: DP1. The first, Data Preview 0 (DP0)[3], contained data products produced from the processing of simulated LSST-like data sets, together with a very early version of the Rubin Science Platform (Jurić et al. 2019). DP1 contains data products derived from the reprocessing of science-grade exposures acquired by the Rubin Commissioning Camera (), in late 2024. The third and final Data Preview, Data Preview 2 (DP2)), is planned to be based on a reprocessing of all science-grade data taken with the Rubin's LSST Science Camera (), during commissioning, and is expected to be released around mid-2026.

All Rubin Data Releases and Previews are subject to a two-year proprietary period, with immediate access granted exclusively to data rights holders (Blum & the Rubin Operations Team 2020). Data rights holders are individuals or institutions with formal authorization to access proprietary data collected by the Vera C. Rubin Observatory. This includes all scientists in the United States, Chile, and designated individuals or groups from other countries[4]. After the two-year proprietary period, DP1 will be made public.

In this paper, we present the contents and validation of, and the data access and community support services for, Rubin DP1, the first Data Preview to deliver data derived from observations conducted by the Vera C. Rubin Observatory. DP1 is based on the reprocessing of a subset of 1792 science-grade exposures acquired over 48 nights during the first on-sky commissioning campaign using the Rubin Commissioning Camera, LSSTComCam, between 2024-11-09 and 2024-12-11. It covers a total area of approximately ∼15 sq. deg. distributed across seven distinct non-contiguous fields. The data products include raw and calibrated single-epoch images, coadded images, difference images, detection catalogs, and other derived data products. DP1 is about 3.5 TB in size and contains around 2.3 million distinct astronomical objects, detected in 2644 coadded images. Full DP1 release documentation is available at https://dp1.lsst.io. Despite Rubin Observatory still being in commissioning and not yet complete, Rubin DP1 provides an important first look at the data, showcasing its characteristics and capabilities.

The structure of this paper is as follows. In §2 we describe the observatory system and overall construction completion status at the time of data acquisition, the seven fields included in DP1 and the observing strategy used. §3 summarizes the contents of DP1 and the different types of data products contained in the release. The data processing pipelines are described in §4, followed by a description of the data validation and performance assessment in §5. §6 describes the Rubin Science Platform (RSP), a cloud-based data science infrastructure that provides tools and services to Rubin data rights holders to access, visualize and analyze peta-scale data generated by the LSST. §7 presents Rubin's model for community support, which emphasizes self-help via documentation and tutorials, and employs an

[2] Shared risk means early access with caveats: the community benefits from getting a head start on science, preparing analyses, and providing feedback, while also accepting that the experience may not be as polished or reliable as it will be during full operations.

[3] See https://dp0.lsst.io

[4] See https://www.lsst.org/scientists/international-drh-list

open platform for asynchronous issue reporting that enables crowd-sourced solutions. Finally, a summary of the DP1 release and information on expected future releases of data is given in §8. The appendix contains a useful glossary of terms and the bibliography.

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

2. ON-SKY COMMISSIONING CAMPAIGN

The first Rubin on-sky commissioning campaign was conducted using the LSSTComCam between 2024-10-24 and 2024-12-11, spanning a total of 48 nights. The primary objective was to optically align the Simonyi Survey Telescope and verify its ability to deliver acceptable image quality using LSSTComCam. In addition, the campaign provided valuable operations experience to facilitate commissioning the full LSSTCam, (Roodman et al. 2024; Lange et al. 2024). It is important to note that commissioning LSSTComCam was not an objective of the campaign. Instead, LSSTComCam was used as a tool to support broader observatory commissioning, including early testing of the Active Optics System (AOS) and the LSST Science Pipelines. As a result, many artifacts present in the data are specific to LSSTComCam and will only be addressed if they persist with LSSTCam. Accordingly, the image quality achieved during this campaign, and in the DP1 data, may not reflect the performance ultimately expected from LSSTCam.

Approximately 16,000 exposures[5] were collected during this campaign, the majority in support of AOS commissioning, system-level verification, and end-to-end testing of the telescope's hardware and software. This included over 10000 exposures for AOS commissioning, more than 2000 bias and dark calibration frames, and over 2000 exposures dedicated to commissioning the LSST Science Pipelines. For DP1, we have selected a subset of 1792 science-grade exposures from this campaign that are most useful for the community to begin preparing for early science.

At the time of the campaign, the observatory was still under construction, with several key components, such as dome thermal control, full mirror control, and the final AOS configuration either incomplete or still undergoing commissioning. As a result, image quality varied widely throughout the campaign and exhibited a broader distribution than is expected with LSSTCam. Despite these limitations, the campaign successfully demonstrated system integration and established a functional observatory.

2.1. *Simonyi Survey Telescope*

The Simonyi Survey Telescope (Stalder et al. 2024) features a unique three-mirror design, including an 8.4-meter Primary Mirror Tertiary Mirror (M1M3) fabricated from a single substrate and a 3.5-meter Secondary Mirror (M2). This compact configuration supports a wide 3.5-degree field of view while enabling exceptional stability, allowing the telescope to slew and settle in under five seconds. To achieve the scientific goals of the 10-year LSST, the Observatory must maintain high image quality across its wide 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 surface deformation primarily under the effect of gravitational and thermal loads.

The AOS, which comprises open- and closed-loop components, optimizes image quality by aligning the camera and M2 relative to M1M3, as well as adjusting the shapes of all three mirrors. The AOS open-loop component corrects for distortions and misalignments resulting from gravitational and thermal effects, while the closed-loop component addresses unpredictable or slowly varying aberrations using feedback from the corner wavefront sensors. The closed-loop wavefront sensing technique is curvature sensing, analyzing extra-focal and intra-focal images to infer the wavefront errors in the system (Thomas et al. 2023). Since LSSTComCam lacks wavefront sensors, wavefront errors were estimated by defocusing the telescope ± 1.5 mm on either side of focus and applying the curvature wavefront pipeline to measure and correct for wavefront errors.

Each night began with an initial alignment correction using a laser tracker to position the system within the capture range of the closed-loop algorithm (Homar et al. 2024). Alignment was achieved using the AOS system. Once the optics were aligned, the image quality was optimized across the LSSTComCam field of view by applying additional corrections to the shape of the mirrors. During Science Pipelines commissioning (§2.4), observations were undertaken using the open-loop component with no correction for thermal effects. The image quality for these data was monitored by measuring the Point Spread Function (PSF) FWHM and periodically rerunning the closed-loop component when the image quality degraded. Under favorable seeing conditions, the delivered image quality was typically around $0.7''$, with a best recorded value of $0.58''$.

[5] We define an "exposure" as the process of exposing all LSSTComCam detectors. It is synonymous with "visit" in DP1. By contrast, an "image" is the output of a single LSSTComCam detector following an exposure.

2.2. *The Rubin Camera*

The LSSTComCam [6], is a 144-megapixel, scaled-down version of the 3.2-gigapixel LSSTCam. It covers approximately 5% of the LSSTCam focal plane area and is designed to validate camera interfaces with other observatory components and evaluate overall system performance prior to the start of LSSTCam commissioning.

The LSSTCam focal plane consists of 21 modular science rafts for imaging, arranged in a 5×5 grid, along with 4 additional corner rafts dedicated to guiding and wavefront sensing. Each raft is a self-contained unit comprising nine 4K×4K Charge-Coupled Device () sensors arranged in a 3×3 mosaic, along with integrated readout electronics and cooling systems. Each sensor is subdivided into 16 segments arranged in a 2×8 layout, with each segment containing 512×2k pixels. All 16 segments are read out in parallel using dedicated amplifiers, one per segment. LSSTCam uses CCD sensors from two vendors: Imaging Technology Laboratory (University of Arizona (UA)) (UA) and Teledyne (E2V). To ensure uniform performance and calibration within each module, individual rafts are populated with sensors from only one vendor.

LSSTComCam consists of a single raft equipped exclusively with ITL sensors. The sensors selected for LSSTComCam represent the lowest-performing units from the LSSTCam production batch and exhibit known issues, including high readout noise (e.g., Detector 8) and elevated Charge Transfer Inefficiency (CTI) (e.g., Detector 5). As a result, some image artifacts observed in the DP1 dataset may be specific to ITL sensors.

Figure 1 shows the single-raft LSSTComCam positioned at the center of the full LSSTCam focal plane. LSSTComCam is designated as Raft 22 (R22) and is installed at the center of the LSSTCam focal plane, corresponding to the central science raft position.

Figure 2 shows the LSSTComCam focal plane layout, illustrating the enumeration of sensors and amplifiers, along with their physical arrangement within the raft. The LSSTCam and LSSTComCam focal planes are described in detail in Plazas Malagón, A. et al. (2025).

LSSTComCam is housed in a support structure that precisely replicates the total mass, center of gravity, and physical dimensions of LSSTCam, with all mechanical and utility interfaces to the telescope implemented identically. This configuration supports full end-to-end testing of the observatory systems, including readout electronics, image acquisition, and data pipelines. The LSSTComCam plate scale is 0.2 arcsec. per pixel.

[6] https://lsstcomcam.lsst.io/

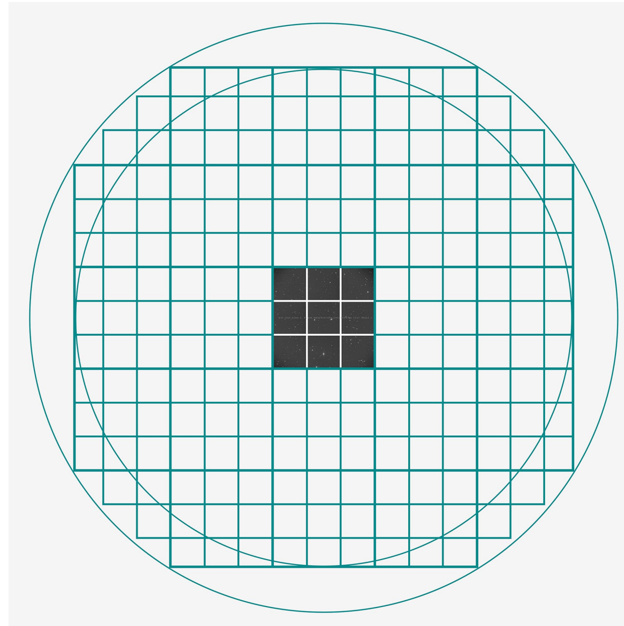

Figure 1. Schematic showing the single-raft LSSTComCam positioned at the center of the full LSSTCam focal plane. The perspective is from above, looking down through the LSSTComCam lenses onto the focal plane. Credit: RubinObs/NOIRLab/SLAC/NSF/DOE/AURA.

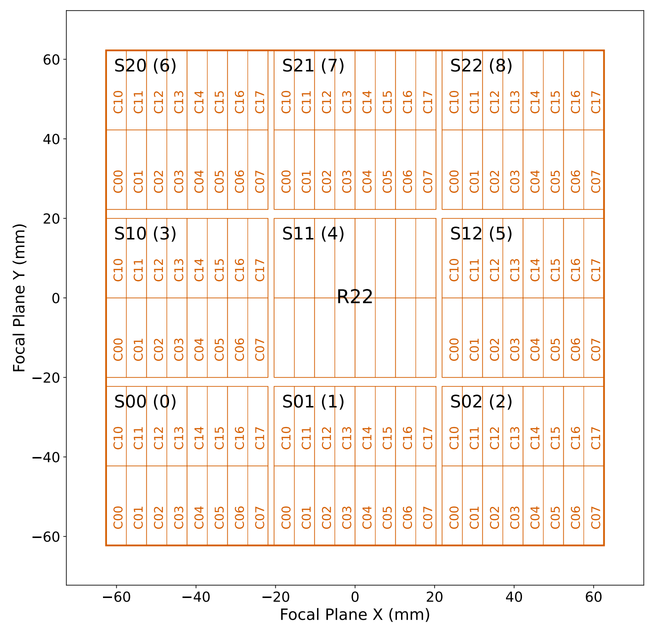

Figure 2. LSSTComCam focal plane layout illustrating the placement and numbering scheme of sensors (S) and amplifiers (C). The view is looking down from above the focal plane through the LSSTComCam lenses. Each sensor contains 16 amplifiers, and a group of nine sensors comprises one raft. LSSTComCam is Raft 22 (R22). The detector number for each sensor is shown in parentheses.

2.2.1. *Filter Complement*

LSSTComCam supports imaging with six broadband filters *ugrizy* spanning 320–1050 nm, identical in design to LSSTCam. However, its filter exchanger can hold only three filters at a time, compared to five in LSSTCam. The full-system throughput of the six LSSTComCam filters, which encompasses contributions from a standard atmosphere at airmass 1.2, telescope optics, camera surfaces, and the mean ITL detector quantum efficiency is shown in Figure 3.

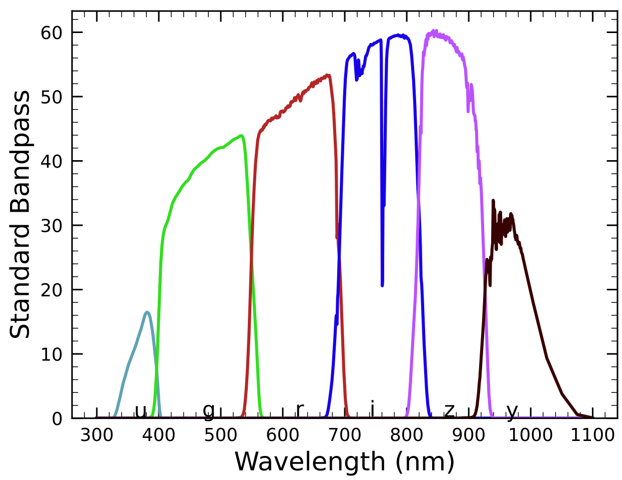

Figure 3. LSSTComCam standard bandpasses, illustrating full system throughput. The bandpasses include a standard atmosphere at airmass 1.2, telescope optics, camera surfaces, and mean ITL detector quantum efficiency.

2.3. *Flat Field System*

During the on-sky campaign, key components of the Rubin calibration system (Ingraham et al. 2022), including the flat field screen, Collimated Beam Projector (), and the Ekspla tunable laser had not yet been installed. As a result, flat fielding for DP1 relied entirely on twilight flats. While twilight flats pose challenges such as non-uniform illumination and star print-through, they were the only available option during LSSTComCam commissioning and for DP1 processing. To mitigate these limitations, dithered, tracked exposures were taken over a broad range of azimuth and rotator angles to construct combined flat calibration frames. Exposure times were dynamically adjusted to reach target signal levels of between 10,000 and 20,000 electrons. Future campaigns will benefit from more stable and uniform flat fielding using the Rubin flat field system, described in Fagrelius & Rykoff (2025).

2.4. *LSST Science Pipelines Commissioning*

Commissioning of the LSST Science Pipelines (Developers 2025) began once the telescope was able to routinely deliver sub-arcsecond image quality. The goals included testing the internal astrometric and photometric calibration across a range of observing conditions, validating the difference image analysis and Prompt Processing (Lim 2022) framework, and accumulating over 200 visits per band to evaluate deep coadded images with integrated exposure times roughly equivalent to those of the planned LSST Wide Fast Deep (WFD) 10-year depth. To support these goals, seven target fields were selected that span a range of stellar densities, overlap with external reference datasets, and collectively span the full breadth of the four primary LSST science themes. These seven fields form the basis of the DP1 dataset. Figure 4 shows the locations of these seven fields on the sky, overlaid on the LSST baseline survey footprint (Jones 2021; Yoachim 2022; Ivezic 2022; The Rubin Observatory Survey Cadence Optimization Committee 2023, 2025), along with sky coverage of both the LSSTCam and LSSTComCam focal planes. Each of the seven target fields was observed repeatedly in multiple bands over many nights. A typical observing epoch on a given target field consisted of 5-20 visits in each of the three loaded filters. Only images taken as 1x30 second exposures have been included in DP1. All images were acquired using the Rubin Feature-Based Scheduler (FBS), version 3.0 (Naghib et al. 2019; Yoachim et al. 2024). Table 1 lists the seven DP1 fields and their pointing centers, and provides a summary of the band coverage in each.

The temporal sampling distribution of observations per band and per night is shown in Figure 5. Gaps in coverage across some bands arise from the fact that LSSTComCam can only accommodate three filters at a time §2.2. As the campaign progressed, the temporal sampling became denser across all fields, reflecting improved efficiency and increased time allocated for science observations. The Extended Chandra Deep Field-South Survey (ECDFS) field received the most consistent and densest temporal sampling. It is important to note that the time sampling in the DP1 dataset differs significantly from what will be seen in the final LSST data.

All fields except for the low ecliptic latitude field, Rubin_SV_38_7, used random translational and rotational dithers within a 0.2 degree radius around the pointing center (Table 1). The rotational dithers were typically applied at the time of filter changes for operational efficiency, with translational dithers of approximately 1 degree applied between individual visits. The Rubin_SV_38_7 field used a different dither pattern to optimize coverage of Solar System Objects and test Solar System Object linking across multiple nights. These observations used a 2 x 2 grid of LSSTComCam pointings to cover an area of about 1.3 degree x 1.3 degrees. The visits cycled between the grid's four pointing centers, using small random dithers to fill chip gaps with the goal of acquiring 3-4 visits per pointing center per band in each observing epoch.

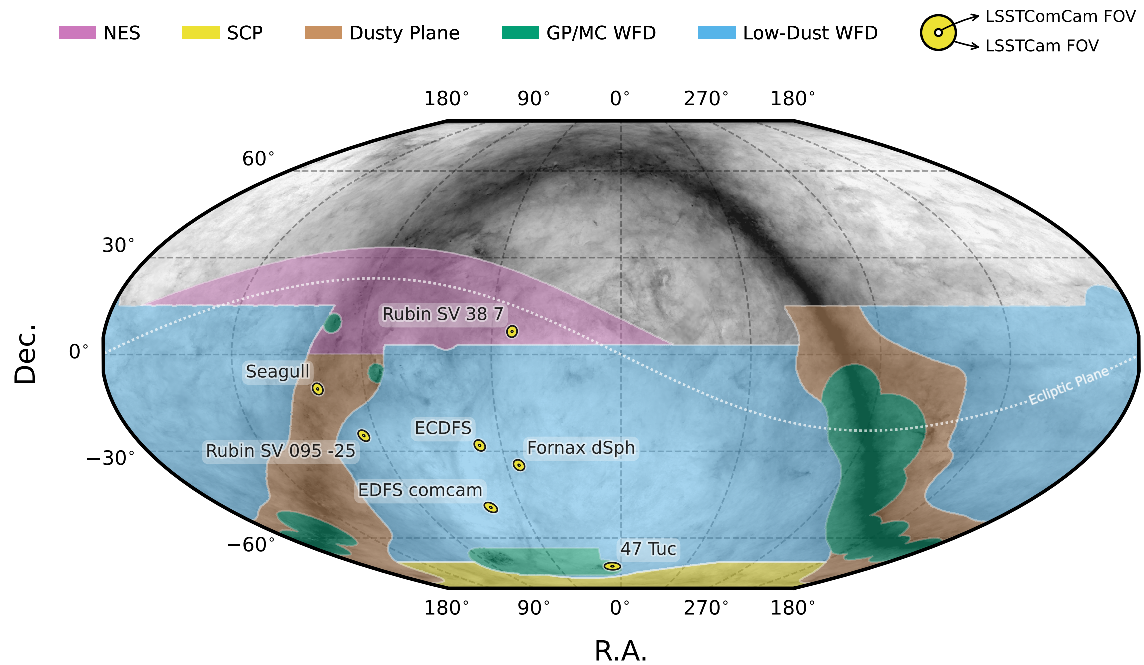

Figure 4. Location of the seven DP1 fields overlaid on the LSST baseline survey footprint. NES: North Ecliptic Spur, SCP: South Celestial Pole, Low-Dust WFD: regions away from the GP observed with a WFD cadence, GP/MC WFD: Galactic Plane and Magellanic Clouds regions observed with a WFD cadence. The field of view (FOV) covered by the LSSTCam and LSSTComCam focal planes is shown as concentric yellow circles about the pointing center of each field.

Table 1. DP1 fields and pointing centers with the number of exposures in each band per field. ICRS coordinates are in units of decimal degrees.

Field Code	Field Name	RA	DEC	Band						Total
		deg	deg	u	g	r	i	z	y	
47_Tuc	47 Tucanae Globular Cluster	6.128	-72.090	6	10	32	19	0	5	72
ECDFS	Extended Chandra Deep Field South	53.160	-28.100	43	230	237	162	153	30	855
EDFS_comcam	Rubin SV Euclid Deep Field South	59.150	-48.730	20	61	87	42	42	20	272
Fornax_dSph	Fornax Dwarf Spheroidal Galaxy	40.080	-34.450	0	5	25	12	0	0	42
Rubin_SV_095_-25	Rubin SV Low Galactic Latitude Field	95.040	-25.000	33	82	84	23	60	10	292
Rubin_SV_38_7	Rubin SV Low Ecliptic Latitude Field	37.980	7.015	0	44	40	55	20	0	159
Seagull	Seagull Nebula	106.300	-10.510	10	37	43	0	10	0	100

2.5. *Delivered Image Quality*

The delivered image quality is influenced by contributions from both the observing system (i.e., dome, telescope and camera) and the atmosphere. During the campaign, the Rubin Differential Image Motion Monitor (DIMM) was not operational, so atmospheric seeing was estimated using live data from the Southern Astrophysical Research Telescope (SOAR) Ring-Image Next Generation Scintillation Sensor () seeing monitor. Although accelerometers mounted on the mirror cell and top-end assembly were available to track dynamic optics effects, such as mirror oscillations that can degrade optical alignment, this data was not used during the campaign. Mount encoder data was used to measure the mount jitter in every image, with a median contribution of 0.004 arcseconds to image degradation measured. As the pointing model was not fine tuned, tracking errors could range from 0.2 to 0.4 arcseconds per image, depending on RA and Dec. Dome and mirror-induced seeing were not measured during the campaign. The median delivered image quality for commanded in-focus images (all bands) was 1.14″, as measured by the PSF FWHM. The best images achieved a PSF FWHM of approximately 0.58″. Ongoing efforts aim to quantify all sources of image degradation, including contributions from the camera system, static and dynamic optical

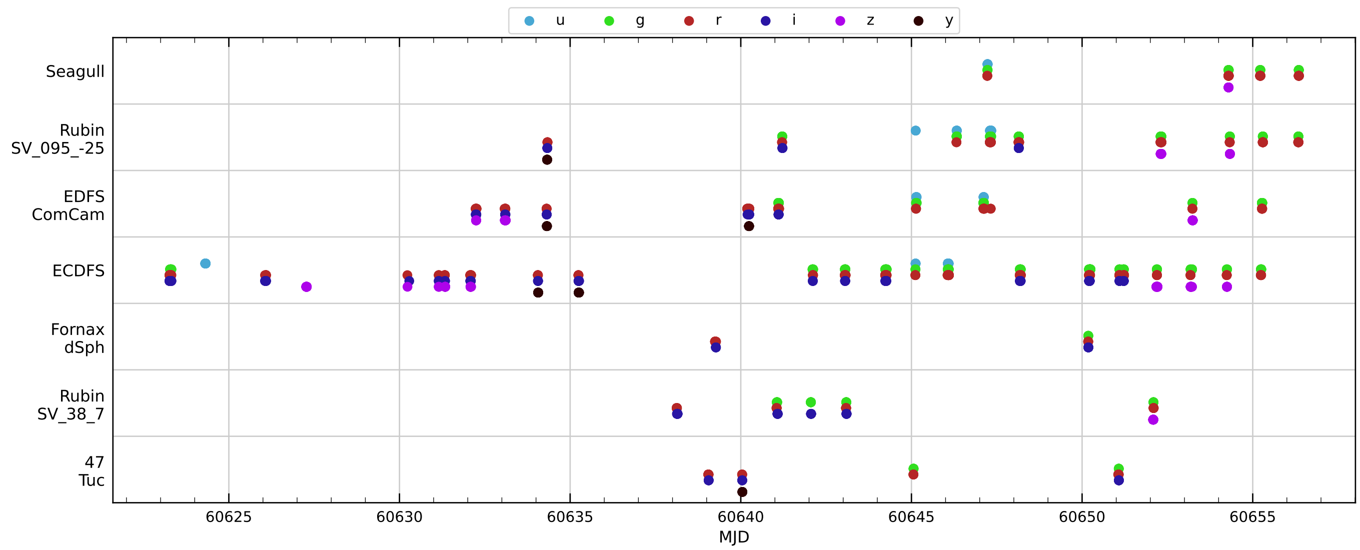

Figure 5. Distribution of DP1 observations by date grouped by field and color coded by band.

components, telescope mount motion, observatory-induced seeing from the dome and mirror, and atmospheric conditions.

3. OVERVIEW OF THE CONTENTS OF RUBIN DP1

Here we describe Rubin DP1 data products and provide summary statistics for each. The DP1 science data products are derived from the 15972 individual CCD images taken across 1792 exposures in the seven LSSTComCam commissioning fields (§2.4).

The data products that comprise DP1 provide an early preview of future LSST data releases and are strongly dependent on the type and quality of the data that was collected during LSSTComCam on-sky campaign (§2.4). Consequently not all anticipated LSST data products, as described in the Data Product Definition Document () (Jurić et al. 2023) were produced for the DP1 dataset.

At the highest level, the DP1 data products fall into one of five types:

- **Images**, including single-epoch images, deep and template coadded images, and difference images;
- **Catalogs** of astrophysical sources and objects detected and measured in the aforementioned images. We also provide the astrometric and photometric reference catalog generated from external sources that was used during processing to generate the DP1 data products;
- **Maps**, which provide non-science-level visualizations of the data within the release. They include, for example, zoomable multi-band images and coverage maps;
- **Ancillary data products**, including, for example, the parameters used to configure the data processing pipelines, log and processing performance files, plots and metrics produced during the data processing steps, and calibration data products (e.g. CTI models, brighter-fatter kernels, etc.);

Table 2. Tract coverage of each DP1 field.

Field Code	Tract ID
ECDFS	5062, 5063, 5064, 4848, 4849
Seagull	7850, 7849, 7610, 7611
Rubin_SV_38_7	10464, 10221, 10222, 10704, 10705, 10463
EDFS_comcam	2393, 2234, 2235, 2394
Rubin_SV_095_-25	5305, 5306, 5525, 5526
47_Tuc	531, 532, 453, 454
Fornax_dSph	4016, 4217, 4218, 4017

- **Metadata** in the form of tables containing information about each visit and processed image, such as pointing, exposure time, and a range of image quality summary statistics.

While images and catalogs are expected to be the primary data products for scientific research, we also recognize the value of providing access to other data types to support investigations and ensure transparency.

To facilitate processing, Rubin DP1 uses a single skymap[7] that covers the entire sky area encompassing the seven DP1 fields. The DP1 skymap divides the entire celestial sphere into 18938 tracts, each covering approximately 2.8 sq. deg.. Each tract is further subdivided into 10×10 equally-sized patches, with each patch covering roughly 0.028 sq. deg. Both tracts and patches overlap with their neighboring regions. Since the LSSTComCam only observed ∼15 sq. deg.

[7] A skymap is a tiling of the celestial sphere, organizing large-scale sky coverage into manageable sections for processing and analysis.

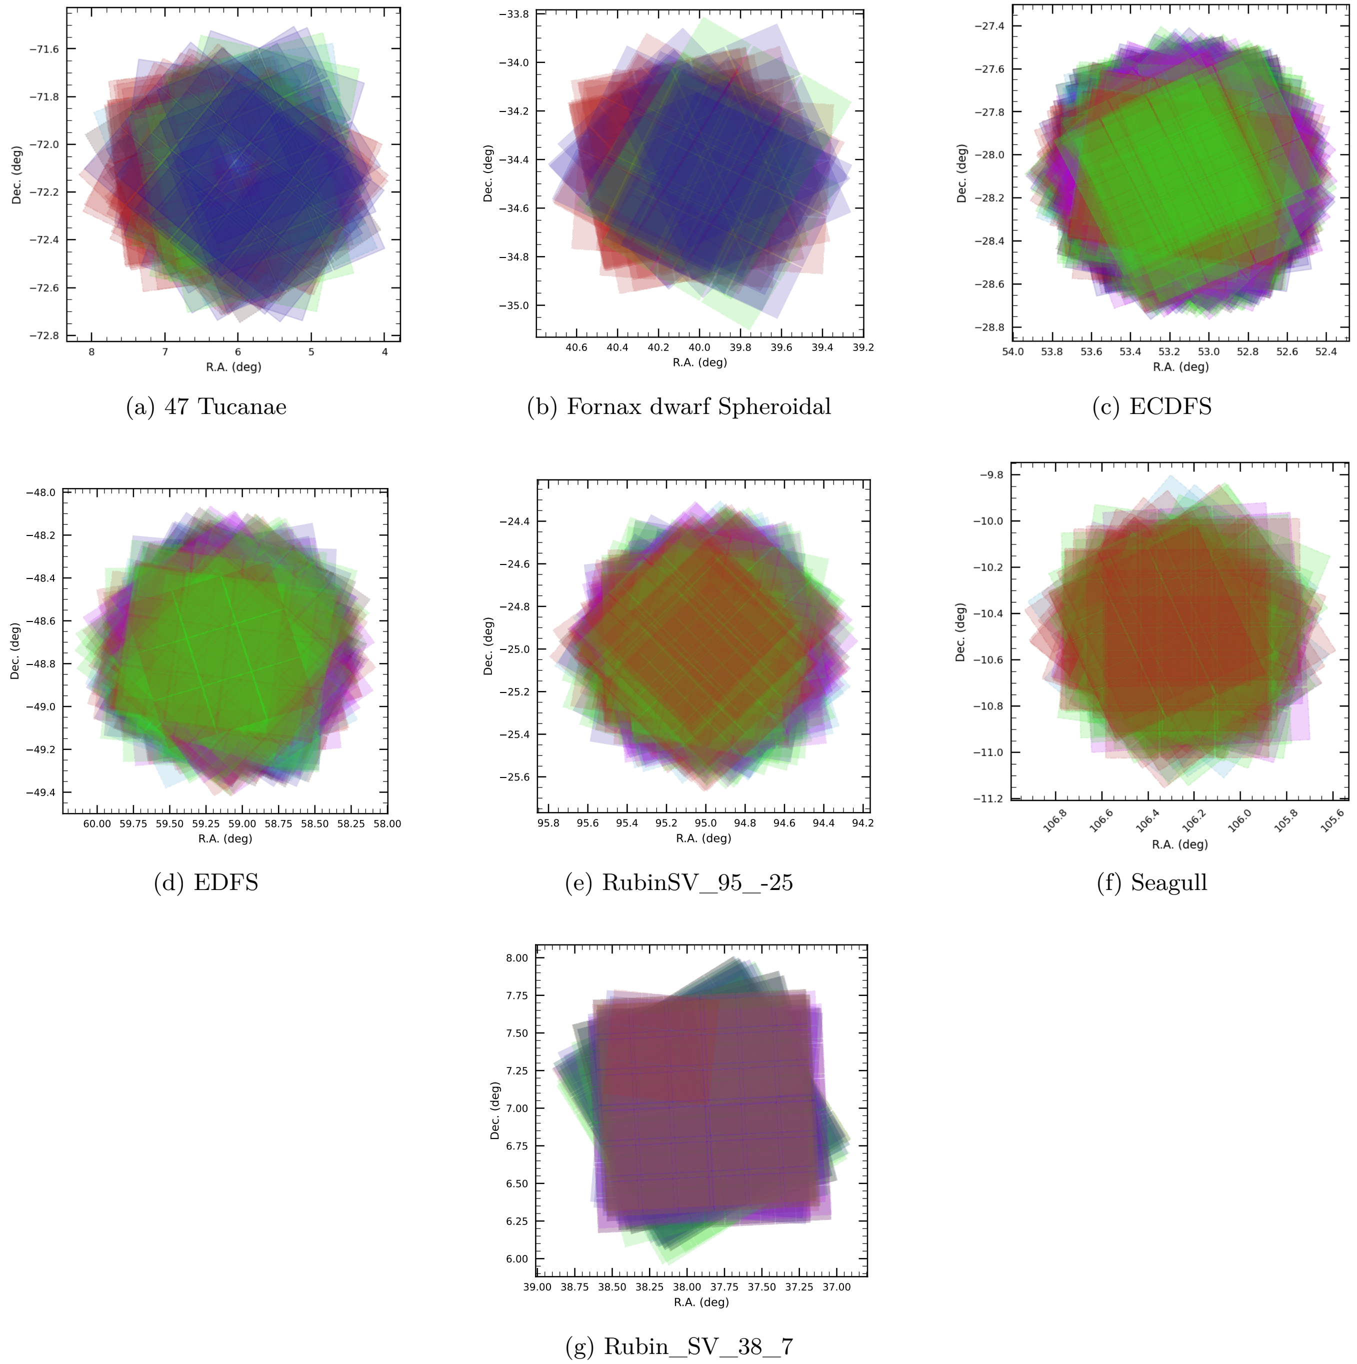

(a) 47 Tucanae (b) Fornax dwarf Spheroidal (c) ECDFS

(d) EDFS (e) RubinSV_95_-25 (f) Seagull

(g) Rubin_SV_38_7

Figure 6. Sky coverage for seven DP1 fields.

of the sky during its campaign, only 29 out of the 18938 tracts have coverage in DP1. The tract identification numbers and corresponding target names for these tracts are listed in Table 2.

The skymap is integral to the production of co-added images. To create a coadded image, the processing pipeline selects all calibrated science images that meet specific quality thresholds (§3.1 and §4.5.1) for a given patch, warps them onto a single consistent pixel grid for that patch, as defined by the skymap, then coadds them. Each individual coadd image therefore covers a single patch. Coadded images and the catalogs of detections from them are termed `tract`-level data products. By contrast, `visit`-level data products are those derived from individual LSSTComCam exposures, such as a raw image or a catalog of detections from a single calibrated image. Most science data products (i.e., images and catalogs) in DP1 are either `tract` or `visit`–level, the main exception being the `Calibration` reference catalog.

Throughout this section, the data product names are indicated using `monospace` font. Data products are accessed

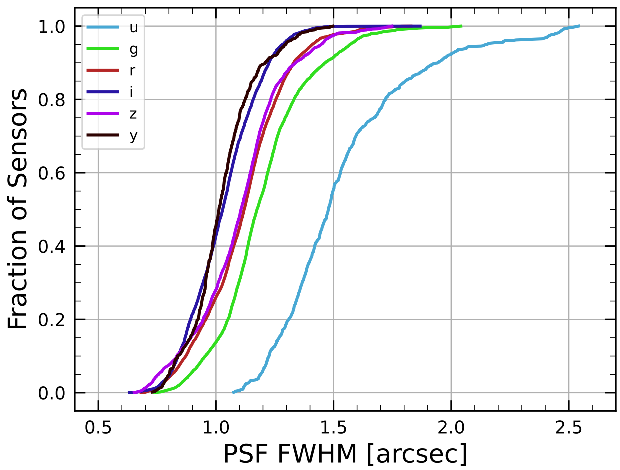

Figure 7. Cumulative distribution of PSF FWHM of the DP1 dataset.

via either the International Virtual-Observatory Alliance (IVOA) Services (§6.2.1) or the Data Butler (§6.2.2), or both.

3.1. *Science Images*

Science images are exposures of the night sky, as distinct from calibration images (§3.5.3). Although the release includes calibration images, allowing users to reprocess the raw images if needed, this is expected to be necessary only in rare cases. Users are strongly encouraged to start from the visit-level images provided. The data product names shown here are those used by the Data Butler, but the names used in the IVOA Services differ only slightly in that they are prepended by "`lsst.`".

- `raw` images are unprocessed data received directly from the camera. Each `raw` corresponds to a single CCD from a single LSSTComCam exposure of 30 s duration. Each LSSTComCam exposure typically produces up to nine `raw`s, one per sensor in the focal plane. However, a small number of exposures resulted in fewer than nine `raw` images due to temporary hardware issues or readout faults.

  In total, DP1 includes 16125 `raw` images. Table 3 provides a summary by target and band. A `raw` contains 4608 × 4096 pixels, including prescan and overscan, and occupies around 18 MB of disk space.[8] The field of view of a single `raw`, excluding prescan and overscan regions, is roughly 0.23° ×0.22° ≈0.051 sq. deg., corresponding to a plate scale of 0.2 arcsec. per pixel.

- `visit_image`s are fully-calibrated processed images. They have undergone instrument signature removal (§4.2.1) and all the single frame processing steps described in §4.2 which are, in summary: PSF modeling, background subtraction, and astrometric and photometric calibration. As with `raw`s, a `visit_image` contains processed data from a single CCD resulting from a single 30 s LSSTComCam exposure. As a consequence, a single LSSTComCam exposure typically results in nine `visit_image`s. The handful of exposures with fewer than nine `raw` images also have fewer than nine `visit_images`, but there are an additional 153 `raw` that failed processing and for which there is thus no corresponding `visit_image`. Almost all failures were due to challenges with astrometric fits or PSF models.

  In total, there are 15972 `visit_image`s in DP1. Each `visit_image` comprises three images: the calibrated science image, a variance image, and a pixel mask, indicating, for example, bad or saturated pixels, pixels affected by cosmic rays, pixels associated with detected sources, etc.). Each `visit_image` also contains a position-dependent PSF model, World Coordinate System () information, and various metadata providing information about the observation and processing. The science and variance images and the pixel mask each contain 4072 × 4000 pixels. In total, a single `visit_image`, including all extensions and metadata, occupies around 110 MB of disk space.

- `deep_coadd`s are the product of warping and co-adding multiple `visit_image`s covering a given patch, as defined by the skymap. `deep_coadd`s are created on a per-band basis, meaning only data from exposures taken with a common filter are coadded. As such, there are up to six `deep_coadd`s covering each patch – one for each of the six LSSTComCam bands. The process of producing `deep_coadd`s is described in detail in §4.5 but, to summarize, it involves the selection of suitable `visit_images` (both in terms of patch coverage, band, and image quality), the warping of those `visit_images` onto a common pixel grid, and the co-adding of the warped `visit_images`. To be included in a DP1 `deep_coadd`, a `visit_image` needed to have a PSF FWHM smaller than 1.7″. Of the 15972 `visit_images`, N satisfied this criterion and were therefore used to create `deep_coadds`.

  There are a total of 2644 `deep_coadds` in DP1. As mentioned above, a single `deep_coadd` covers one patch, and includes a small amount of overlap with its neighboring patch. The skymap used for DP1 defines a patch as having an on-sky area of 0.028 sq. deg. excluding overlap, and 0.036 sq. deg. including overlap. A single `deep_coadd` – including overlap – contains 3400 × 3400 equal-sized pixels, corresponding to a platescale of 0.2 arcsec. per pixel. Each `deep_coadd` contains the science image (i.e., the coadd), a variance image, and a pixel mask; all three contain the same number of pixels. Each `deep_coadd` also contains a position-dependent PSF model (which is the weighted sum of the PSF models of the input `visit_images`), WCS information, plus various metadata.

[8] Each amplifier image contains 3 and 64 columns of serial prescan and overscan pixels, respectively, and 48 rows of parallel overscan pixels, meaning a `raw` contains 4072 ×4000 exposed pixels.

Table 3. Number of `raw` per field and band.

Field Code	Band						Total
	u	g	r	i	z	y	
47_Tuc	54	90	288	171	0	45	648
ECDFS	387	2070	2133	1455	1377	270	7692
EDFS_comcam	180	549	783	378	378	180	2448
Fornax_dSph	0	45	225	108	0	0	378
Rubin_SV_095_-25	297	738	756	207	540	90	2628
Rubin_SV_38_7	0	396	360	495	180	0	1431
Seagull	90	333	387	0	90	0	900
Total	1008	4221	4932	2814	2565	585	16125

Since coadds always cover an entire patch, it is fairly common for a deep_coadd to contain regions that were not covered by any of the selected `visit_image`s, particularly if the patch is on the outskirts of a field and was thus not fully observed. By the nature of how coadds are produced, such regions may contain seemingly valid flux values (i.e., not necessarily zeros or `NaN`s), but will instead be flagged with the `NO_DATA` flag in the pixel mask. It is therefore crucial that the pixel mask is referred to when analyzing `deep_coadds`.

- `template_coadd`s are those created to use as templates for difference imaging, i.e., the process of subtracting a template image from a `visit_image` to identify either variable or transient objects.[9] As with `deep_coadds`, `template_coadds` are produced by warping and co-adding multiple `visit_image`s covering a given skymap-defined patch. The process of building `template_coadd`s is the same as that for `deep_coadds`, but the selection criteria differ between the two types of coadd. In the case of `template_coadds`, the third of `visit_image`s covering the patch in question with the smallest PSF FWHM are selected. If one third corresponds to fewer than twelve `visit_images` (i.e., there are fewer than 36 `visit_image`s covering the patch), then the twelve `visit_images` with the smallest PSF FWHM are selected. Finally, if there are fewer than twelve `visit_images` covering the patch, then all `visit_image`s are selected. Of the 15972 `visit_images`, N were used to create `template_coadd`s. This selection strategy is designed to optimize for seeing when a patch is well-covered by `visit_image`s, yet still enabling the production of `template_coadd`s for poorly-covered patches.

  There are a total of 2730 `template_coadd`s in DP1.[10] As with `deep_coadd`s, a single `template_coadd` covers a single patch. Since the same `skymap` is used when creating both `deep_coadd` and `template_coadds`, the on-sky area and pixel count of `template_coadd`s are the same as that of a `deep_coadd` (see above). Similarly, `template_coadd`s contain the science image (i.e., the coadd), a variance image, and a pixel mask; all three contain the same number of pixels. Also included is the PSF model, WCS information, and metadata. As is the case for `deep_coadd`, those pixels within `template_coadd`s that are not covered by any of the selected `visit_image`s may still have seemingly valid values, but are indicated with the `NO_DATA` flag within the pixel mask.

- `difference_image`s are generated by the subtraction of the warped, scaled, and PSF-matched `template_coadd` from the `visit_image` (see §4.6.1). In principle, only those sources whose flux has changed relative to the `template_coadd` should be apparent (at a significant level) within a `difference_image`. In practice, however, there are numerous spurious sources present in `difference_image`s due to unavoidably imperfect template matching.

  In total, there are 15972 `difference_image`s in DP1, one for each `visit_image`.

  Like `visit_image`s, `difference_image`s contain the science (i.e., difference) image, a variance image, and a pixel mask; all three contain the same number of pixels, which is the same as that of the input `visit_image`. Also included is the PSF model, WCS information, and metadata.

- Background images contain the model background that has been generated and removed from a science image. `visit_image`s, `deep_coadds` and `template_-`

[9] It should be noted that `template_coadd`s are not themselves subtracted from `visit_image`s but are, instead, warped to match the WCS of a `visit_image`. It is this warped template that is subtracted from the `visit_image` to create a difference image. For storage space reasons, warped templates are not retained for DP1, as they can be readily and reliably recreated from the `template_coadd`s.

[10] The difference in the number of `deep_coadd`s and `template_coadd`s is due to the difference in the `visit_image` selection criteria for each coadd.

coadds all have associated background images.[11] Background images contain the same number of pixels as their respective science image, and there is one background image for each visit_image, deep_coadd, and template_coadd. Difference imaging analysis also measures and subtracts a background model, but the difference_background data product is not written out by default and is not part of DP1.

Background images are not available via the IVOA Service; they can only be accessed via the Butler Data Service.

3.2. *Catalogs*

Here we describe science-ready tables produced by the science pipelines. All but one of the catalogs described here contain data for detections in the images described in §3.1, the exception being the Calibration catalog that contains reference data obtained from previous surveys. Observatory-produced metadata tables are described in §3.4

The Rubin Observatory has adopted the convention by which single-epoch detections are referred to as sources. By contrast, the astrophysical object associated with a given detection is referred to as an object. [12] As such, a given object will likely have multiple associated sources, since it will be observed in multiple epochs. Each type of catalog contains measurements for either sources or objects detected in one of visit_images, deep_coadds, or difference_images.

While the Source, Object, ForcedSource, DiaSource, DiaObject, and ForcedSourceOnDiaObject catalogs described below each differ in terms of their specific columns, in general they each contain: one or more unique identification number, positional information, one or more types of flux measurements (e.g., aperture fluxes, PSF fluxes, Gaussian fluxes, etc.), and a series of boolean flags (indicating, for example, whether the source/object is affected by saturated pixels, cosmic rays, etc.) for each source/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, all measured properties are reported with their associated 1σ uncertainties.

Since DP1 is a preview, it doesn't include all the catalogs expected in a full LSST Data Release. Additionally, the catalogs it does include may be missing some columns planned for future releases. Where this is the case, we note what data is missing in the catalog descriptions that follow.

Catalog data are stored in the Qserv database (§6.5.1) and are accessible via Table Access Protocol (IVOA standard) (IVOA), and an online DP1 catalog schema is available at https://sdm-schemas.lsst.io/dp1.html. Catalog data are also accessible via the Data Butler (§6.2.2).

[11] In future data releases, background images may be included as part of their respective science image data product.

[12] We caution that this nomenclature is not universal; for example, some surveys call "detections" what we call "sources", and use the term "sources" for what we call "objects".

- The Source catalog contains data on all sources which are, prior to deblending, detected with a greater than 5σ significance in each individual visit. The detections reported in the Source catalog have undergone deblending; in the case of blended detections, only the deblended sources are included in the Source catalog. It is important to note that while the criterion for inclusion in a Source catalog is a $> 5\sigma$ detection in a visit_image prior to deblending, the positions and fluxes are reported post-deblending. Hence, it is possible for the Source catalog to contain sources whose flux-to-error ratios – potentially of all types (i.e., aperture flux, PSF flux, etc.) – are less than 5.

  In addition to the general information mentioned above (i.e., IDs, positions, fluxes, flags), the Source catalog also include basic shape and extendedness information.

  The Source catalog contains data for 46 million sources in DP1.

- 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 cross-matching has been performed between bands to associate an object in one band with its counterpart(s) in the other bands. As such, unlike the Source catalog, the Object catalog contains data from 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 post-deblending. Hence, it is possible for Object catalog to contain objects whose flux-to-error ratios — potentially of all types and in all bands — are less than 5.

  In addition to the general information mentioned above (i.e., IDs, positions, fluxes, 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 is included in the DP1 object catalogs.

  The Object catalog contains data for 2.3 million objects in DP1.

- The ForcedSource catalog contains forced PSF photometry measurements performed on both difference_images (i.e., the psfDiffFlux column) and visit_images (i.e., the psfFlux column) at the positions of all the objects in the Object catalog. We recommend using the psfDiffFlux column when generating lightcurves because they are less sensitive to flux from neighboring sources. As well as forced photometry PSF fluxes, a range of boolean flags are also included in the `ForcedSource` catalog.

The `ForcedSource` catalog contains a total of 269 million entries across 2.3 million unique objects.

- The `DiaSource` catalogs contains data on all the sources detected at a $> 5\sigma$ significance — including those associated with known Solar System objects — in the `difference_images`. Unlike sources detected in `visit_image`, sources detected in difference images (hereafter, "DiaSources") have gone through an association step during which an attempt has been made to associate them with into underlying objects called "DiaObject"s. The `DiaSource` catalog consolidates all this information across multiple visits and bands. The detections reported in the `DiaSource` catalog have not undergone deblending.

  The `DiaSource` catalog contains data for 3.1 million `DiaSources` in DP1.

- The `DiaObject` catalog contains the astrophysical objects that DiaSources are associated with (i.e., the "DiaObjects"). The `DiaObject` catalog only contains non-Solar System Objects; Solar System Objects are, instead, recorded in the `SSObject` catalog (see below for a description of the `SSObject` catalog). When a DiaSource is identified, the `DiaObject` and `SSObject` catalogs are searched for objects to associate it with. If no association is found, a new DiaObject is created and the DiaSource is associated to it. Along similar lines, an attempt has been made to associate DiaObjects across multiple bands, meaning the `DiaObject` catalog – like the `Object` catalog – contains data from multiple bands. Since DiaObjects are typically transient or variable (by the nature of their means of detection), the `DiaObject` catalog contains summary statistics of their fluxes, such as the mean and standard deviation over multiple epochs; users must refer to the `ForcedSourceOnDiaObject` catalog (see below) or the `DisSource` catalog for single epoch flux measurements of DiaObjects.

  The `DIAObject` catalogs contains data for 1.1 million DiaObjects in DP1.

- The `ForcedSourceOnDiaObject` catalog is equivalent to the `ForcedSource` catalog, but contains forced photometry measurements obtained at the positions of all the DiaObjects in the `DiaObject` catalog.

  The `ForcedSourceOnDiaObject` catalog contains a total of 197 million entries across 1.1 million unique DiaObjects.

- The `CcdVisit` catalog contains data for each individual processed `visit_image`. In addition to technical information, such as the on-sky coordinates of the central pixel and measured pixel scale, the `CcdVisit` catalog contains a range of data quality measurements, such as whole-image summary statistics for the PSF size, zeropoint, sky background, sky noise, quality of astrometry solution. It provides an efficient method to access `visit_image` properties without needing to access the image data.

  The `CcdVisit` catalog contains entries summarizing data for all 16071 `visit_images`.

- The `SSObject` catalog..... Two tables, named Minor Planet Center Orbit database () and SSObject, carry information about Solar System Objects. The MPCORB table provides the Minor Planet Center-computed orbital elements for all known asteroids, including Rubin-discovered. For DP1, the SSObject serves primarily to provide the mapping between the International Astronomical Union (IAU) designation of an object (listed in MPCORB), and the internal ssObjectId identifier (used as a key to find solar system object observations in the DiaSource and SSSource tables).

- The `SSSource` catalog contain data on all DiaSources that are either associated with previously-known Solar System Objects, or have been confirmed as newly-discovered Solar System Objects by confirmation of their orbital properties. As entries in the `SSSource` catalog stem from the `DiaSource` catalog, they have all been detected at $> 5\sigma$ significance.

  The `SSSource` catalog contains data for 5988 Solar System Sources.

- The `Calibration` catalog is the reference catalog that was used to perform astrometric and photometric calibration. It is a whole-sky catalog built specifically for LSST, as no prior reference catalog had both the depth and coverage needed to calibrate LSST data. It combines data from multiple previous reference catalogs and contains only stellar sources. Full details on how the `Calibration` catalog was built are provided in Ferguson et al. (2025) (in which it is referred to as "The Monster"), but we provide a brief summary here.

  For the *grizy* bands, the input catalogs were (in order of priority): Dark Energy Survey (DES) Y6 Calibration Stars (Rykoff et al. 2023); Gaia-B or R Photometry (Gaia) () Synthetic Magnitudes (Gaia Collaboration et al. 2023); the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS)1 3pi survey (Chambers et al. 2016); Data Release 2 of the the SkyMapper survey (Onken et al. 2019); and Data Release 4 of the Very Large Telescope (European Southern Observatory (ESO)) (ESO) Survey Telescope (ESO) Asteroid Terrestrial-impact Last Alert System () survey (Shanks et al. 2015). For the *u*-band, the input catalogs were (in order of priority): Standard Stars from Sloan Digital Sky Survey () 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. This latter input (i.e., SLR estimates) was 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 large footprint of the LSST.

  Only high quality stellar sources were selected from each input catalog. Throughout, the `Calibration` 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/SDSS bandpasses) to the respective bandpasses. All sources from the input catalogs are matched to *Gaia*-Data Release 3 () sources for robust astrometric information, selecting only isolated sources (i.e., no neighbors within 1″).

Once the input catalogs had been collated and fluxes transformed to the standard DES/SDSS bandpasses, the LSST Science Pipeline's ConvertReferenceCatalogTask was used to shard the catalog, which allows it to be quickly searched for sources covering a particular patch of sky, and create a set of standard columns containing positional and flux information, including uncertainties.

3.3. *Maps*

Maps are two-dimensional visualizations of survey data. In DP1, these fall into two categories: Survey Property Maps and Hierarchical Progressive Survey (IVOA standard) (IVOA) Maps (Fernique et al. 2015).

3.3.1. *Survey Property Maps*

Survey Property Maps summarize how properties such as observing conditions or exposure time vary across the observed sky. Each map provides the spatial distribution of a specific quantity at a defined sky position for each band by aggregating information from the images used to make the deep_coadd. Maps are initially created per-tract and then combined to produce a final consolidated map. At each sky location, represented by a spatial pixel in the Hierarchical Equal-Area iso-Latitude Pixelisation (HEALPix) grid, values are derived using statistical operations, such as minimum, maximum, mean, weighted mean, or sum, depending on the property.

There are 29 survey property maps in DP1. The available maps describe total exposure times, observation epochs, PSF size and shape, PSF magnitude limits, sky background and noise levels, as well as astrometric shifts and PSF distortions due to wavelength-dependent atmospheric Differential Chromatic Refraction () effects. They all use the dataset type format deep_coadd_<PROPERTY>_consolidated_map_<STATISTIC> e.g. deep_coadd_exposure_time_consolidated_map_sum provides a spatial map of the total exposure time accumulated per sky position in units of seconds. All maps are stored in HealSparse[13](Górski et al. 2005) format. Survey property maps are only available via the Data Butler (§6.2.2) and have dimensions band and skymap.

Figure 8 presents three survey property maps for exposure time, PSF magnitude limit, and sky noise, computed for representative tracts and bands. Because full consolidated maps cover widely separated tracts, we use clipped per-tract views here to make the spatial patterns more discernible.

3.3.2. *HiPS Maps*

HiPS Maps (Fernique et al. 2015), offer an interactive way to explore seamless, multi-band tiles of the sky regions covered by DP1, allowing for smooth panning and zooming. DP1 provides multi-band HiPS images created by combining data from individual bands of deep_coadd and template_coadd images. These images are false-color representations generated using various filter combinations for the red, green, and blue channels. The available filter combinations include *gri*, *izy*, *riz*, and *ugr* for both deep_coadd and template_coadd. Additionally, for deep_coadd only, we provide color blends such as *uug* and *grz*. Post-DP1, we plan to also provide single-band HiPS images for all *ugrizy* bands in both Portable Network Graphics (PNG) and Flexible Image Transport System () formats.

HiPS maps are only accessible through the HiPS viewer in the Rubin Science Platform () Portal (§6.3) and cannot be accessed via the Data Butler (§6.2.2). All multi-band HiPS images are provided in PNG format.

3.4. *Metadata*

DP1 also includes metadata about the observations, which is stored in the Visit table. The data it contains is produced by the observatory directly, rather than the science pipelines. It contains technical data for each visit, such as telescope pointing, camera rotation, airmass, exposure start and end time, and total exposure time.

3.5. *Ancillary Data Products*

DP1 also includes several ancillary data products. While we do not expect most users to need these, we describe them here for completeness. All the Data Products described in this section can only be accessed via the Data Butler (§6.2.2).

3.5.1. *Task configuration, log, and metadata*

DP1 includes provenance-related data products such as task logs, configuration files, and task metadata. Configuration files record the parameters used in each processing task, while logs and metadata contain information output during processing. These products help users understand the processing setup and investigate potential processing failures.

3.5.2. *Pipeline-generated plots and metrics*

DP1 includes various plots and metrics generated during data processing, such as plots comparing measured fluxes and source positions relative to references, and metrics indicating the numbers of flagged pixels in a given visit_image. These data products are predominantly used by the data management team to assess the quality of the processed data. We include them with DP1 for transparency.

3.5.3. *Calibration Data Products*

Calibration data products include a variety of images and models that are used to characterize and correct the performance of the camera and other system components. These

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

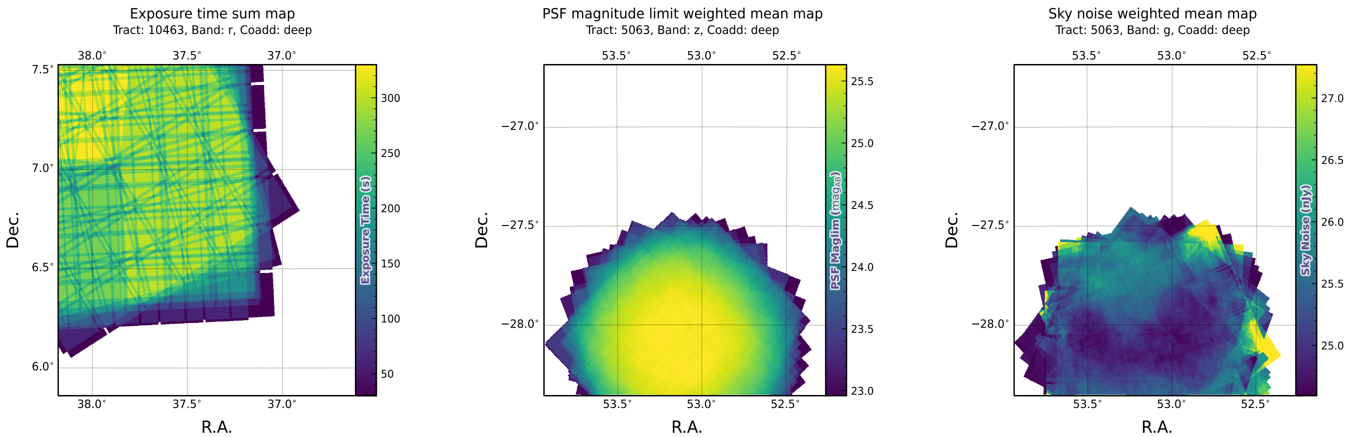

(a) Total exposure time sum map for `deep_coadd` tract 10463, band: r in field Rubin_SV_38_7

(b) 5σ PSF magnitude limit weighted mean map for `deep_coadd` tract 5063, band z in field ECDFS

(c) Sky noise weighted mean map for `deep_coadd` tract 5063, band z in ield ECDFS

Figure 8. Examples of survey property maps from Rubin DP1 across different bands, clipped to the boundary of a single tract for visual clarity.

include bias, dark, and flat-field images, Photon Transfer Curve (PTC) gains, brighter-fatter kernels, charge transfer inefficiency (CTI) models, linearizers, and illumination corrections. For flat-field corrections, DP1 processing used combined flats, which are averaged from multiple individual flat-field exposures to provide a stable calibration. These calibration products are essential inputs to Instrument Signal Removal (ISR) (§4.2.1). While these products are included in DP1 for transparency and completeness, users should not need to rerun ISR for their science and are advised to start with the processed `visit_image`.

3.5.4. *Standard Bandpasses*

The `standard_passband` data products contain the system throughputs described in §2.2.1.

4. DATA RELEASE PROCESSING

Data Release Processing () is the systematic reprocessing of all Rubin Observatory data collected up to a certain date to produce the calibrated images, catalogs of detections, and derived data products described in Section 3. DP1 was processed entirely at the United States Data Facility (USDF), using 17,024 CPU hours.[14]

This section describes the pipeline algorithms used to produce DP1 and how they differ from those planned for full-scale LSST data releases. Data Release Production consists of four major stages: (1) single-frame processing, (2) calibration, (3) coaddition, and (4) difference imaging analysis (Difference Image Analysis ()).

4.1. *LSST Science Pipelines Software*

The LSST Science Pipelines software (Developers 2025; Swinbank et al. 2020) will be used to generate all Rubin Observatory and LSST data products. It provides both the algorithms and middleware frameworks necessary to process raw data into science-ready products, enabling analysis by the Rubin scientific community. Version v29.1 of the pipelines was used to produce DP1. Documentation for this version is available at: https://pipelines.lsst.io/v/v29_1

4.2. *Single Frame Processing*

4.2.1. *Instrument Signature Removal*

The first step in processing LSSTComCam images is to correct for the effects introduced by the telescope and detector. Each sensor and its readout amplifiers can vary slightly in performance, causing images of even a uniformly illuminated focal plane to exhibit discontinuities and shifts due to detector effects. The ISR pipeline aims to recover the original astrophysical signal as best as possible and produce science-ready single-epoch images for source detection and measurement (see Fagrelius & Rykoff 2025; Plazas Malagón et al. 2025 for a detailed description of the ISR procedures).

Figure 9 illustrates the model of detector components and their impact on the signal, tracing the process from photons incident on the detector surface to the final quantized values recorded in the image files. The ISR pipeline essentially "works backward" through the signal chain, correcting the integer analog-to-digital units (ADU) raw camera output back to a floating-point number of photoelectrons created in the silicon. The physical detector, shown on the left in Figure 9, is the source of effects that arise from the silicon itself, such as the dark current and the brighter-fatter effect (Plazas et al. 2018; Broughton et al. 2024). After the image has integrated, the charge is shifted to the serial register and read out, which can introduce charge transfer inefficiencies and a clock-injected offset level. The signals for all amplifiers are transferred via cables to the Readout Electronics Board (REB), during which crosstalk between the amplifiers may occur. The Analog Signal Processing Integrated Circuit (ASPIC) on the REB converts the analog signal from

[14] For future Data Releases, data processing will be distributed across the USDF, the French Data Facility (FrDF) and the United Kingdom Data Facility (UKDF).

the detector into a digital signal, adding both quantization and a bias level to the image. Although the signal chain is designed to be stable and linear, the presence of numerous sources of non-linearity indicates otherwise.

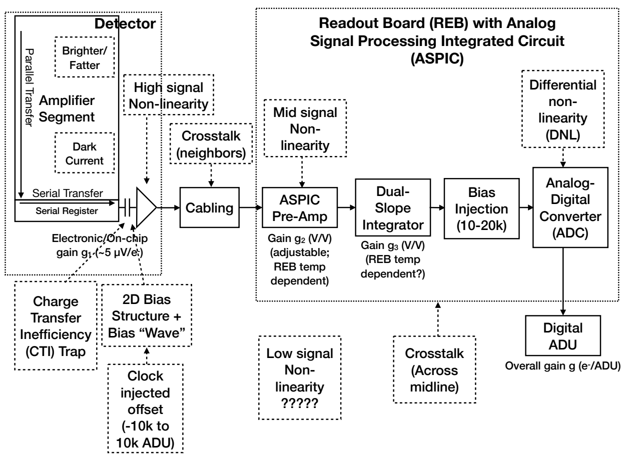

Figure 9. The model of the detector and REB components, labeled with the effects that they impart on signal.

The ISR processing pipeline for DP1 performs, in the following order: Analogue-to-Digital Unit (ADU) dithering to reduce quantization effects, serial overscan subtraction, saturation masking, gain normalization, crosstalk correction, parallel overscan subtraction, linearity correction, serial CTI correction, image assembly, bias subtraction, dark subtraction, brighter-fatter correction, defect masking and interpolation, variance plane construction, flat fielding, and amplifier offset (amp-offset) correction[15]. Flat fielding for DP1 was performed using combined flats produced from twilight flats acquired with sufficient rotational dithering to mitigate artifacts from print-through stars, as described in §2.3.

4.2.2. *Background Subtraction*

The background subtraction algorithms in the LSST Science Pipelines estimate and remove large-scale background signals from science imaging. Such signals may include sky brightness from airglow, moonlight, scattered light instrumental effects and diffuse astrophysical emission. In so doing, true astrophysical sources are isolated to allow for accurate detection and measurement.

To generate a background model, each post-ISR image is divided into superpixels of 128×128 pixels. Pixels with a mask flag set that indicates that they contain no useful science data or that they contain flux from a preliminary source detection are masked. The iterative 3σ clipped mean of the remaining pixels is calculated for each superpixel, constructing a background statistics image. A sixth-order Chebyshev polynomial is fit to these values to allow for an extrapolation back to the native pixel resolution of the post-ISR image.

4.3. *Calibration*

Stars are detected in each post-ISR image using a 5σ threshold. Detections of the same star across multiple images are then associated to identify a consistent set of isolated stars with repeated observations suitable for use in PSF modeling, photometric calibration, and astrometric calibration.

Initial astrometric and photometric solutions are derived using only the calibration reference catalogs (see §3.2), and an initial PSF model is fit using PSFEx (Bertin 2011). These preliminary solutions provide approximate source positions, fluxes, and PSF shapes that serve as essential inputs to the calibration process, enabling reliable source matching, selection of high-quality stars, and iterative refinement of the final astrometric, photometric, and PSF models. These preliminary solutions are subsequently replaced by more accurate fits, as described in the following sections.

4.3.1. *PSF Modeling*

PSF modeling in DP1 uses the Piff (Jarvis et al. 2021) algorithm. Piff models represent the PSF on a pixel-by-pixel basis and interpolate its parameters across a single CCD using two-dimensional polynomials. Piff utilizes its Pixel grid model with a fourth-order polynomial interpolation per CCD, except in the u-band, where star counts are insufficient to support a fourth-order fit. In this case, a second-order polynomial is used instead. Details on the choice of polynomial order, overall PSF modeling performance, and known issues are discussed in §5.2.

4.3.2. *Astrometric Calibration*

Starting from the astrometric solution calculated in single frame processing §4.2, the final astrometric solution is computed using the ensemble of visits in a given band that overlap a given tract. This allows the astrometric solution to be further refined by using all of the isolated point sources of sufficient signal-to-noise ratio in an image, rather than only those that appear in the reference catalog (as is done in single frame processing). Using multiple whole visits rather than a single detector also allows us to account for effects that impact the full focal plane and for the proper motion and parallax of the sources.

In order to perform the fit of the astrometric solution, isolated point sources are associated between overlapping visits and with the Gaia DR3 reference catalog where possible. The model used for DP1 consists of a static map from pixel-space to an intermediate frame (the per-detector model), followed by a per-visit map from the intermediate frame to the plane tangent to the telescope boresight (the per-visit model), then finally a deterministic mapping from the tangent plane to the sky. The fit is done using the `gbdes` package (Bernstein et al. 2017), and a full description is given in Saunders (2024).

The per-detector model is intended to capture quasi-static characteristics of the telescope and camera. During Rubin Operations, the astrometric solution will allow for separate epochs with different per-detector models, to account

[15] Amp-offset corrections are designed to address systematic discontinuities in background sky levels across amplifier boundaries. The implementation in the LSST Science Pipelines is based on the `Pan-STARRS` Pattern Continuity algorithm (Waters et al. 2020)

for changes in the camera due to warming and cooling and other discrete events. However, for DP1, LSSTComCam was assumed to be stable enough that all visits use the same per-detector model. The model itself is a separate two-dimensional polynomial for each detector. For DP1, a degree 4 polynomial was used; the degree of the polynomial mapping is tuned for each instrument and may be different for LSSTCam. Further improvements may be made by including a pixel-based astrometric offset mapping, which would be fit from the ensemble of astrometric residuals, but this is not included in the DP1 processing.

The per-visit model attempts to account for time-varying effects on the path of a photon from both atmospheric sources and those dependent on the telescope position. This model is also a polynomial mapping, in this case a degree 6 two-dimensional polynomial. Correction for DCR was not done for DP1, but will be included in LSSTCam processing during Operations. Future processing will also likely include a Gaussian Processes fit to better account for atmospheric turbulence, as was demonstrated in Fortino et al. (2021) and Léget et al. (2021).

The last component of the astrometric calibration is the position of the isolated point sources included in the fit. The positions consist of five parameters: position on the sky, proper motion, and parallax. The reference epoch for the fit positions is 2024.9.

4.3.3. *Photometric Calibration*

Photometric calibration of the DP1 dataset is based on the Forward Global Calibration Method (FGCM Burke et al. 2018), adapted for the LSST Science Pipelines (Aihara et al. 2022; Fagrelius & Rykoff 2025). We used Forward Global Calibration Model (FGCM) to calibrate the full DP1 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. However, because full scans of the LSSTComCam as-built filters and individual detectors were not available, we instead used the nominal reference throughputs for the Simonyi Survey Telescope and LSSTCam.[16] These nominal throughputs were sufficient for the DP1 calibration, given the small and homogeneous focal plane consisting of only 9 ITL detectors. The FGCM atmosphere model, provided by MODTRAN (Berk et al. 1999), was used to generate a look-up table for atmospheric throughput as a function of zenith distance at Cerro Pachón. This model accounts for Rayleigh scattering by molecular oxygen (O_2) and ozone (O_3), absorption by water vapor, and Mie scattering by airborne aerosol particulates. Nightly variations in the atmosphere are modeled by minimizing the variance in repeated observations of stars with a Signal to Noise Ratio (SNR) greater than 10, measured using "compensated aperture fluxes". These fluxes include a local background subtraction (see §4.2.2 to mitigate the impact of background offsets. The model fitting process incorporates all 6 bands (*ugrizy*) but does not include any gray (achromatic) terms, except for a linear assumption of mirror reflectance degradation, which is minimal over the short duration of the DP1 observation campaign. As an additional constraint on the fit, we use a subset of stars from the reference catalog (Ferguson et al. 2025), primarily to constrain the system's overall throughput and establish the "absolute" calibration.

[16] Available at: https://github.com/lsst/throughputs/tree/1.9

4.4. *Visit Images and Source Catalogs*

With the final PSF models, WCS solutions, and photometric calibrations in place, we reprocess each single-epoch image to produce a final set of calibrated visit images and source catalogs. Source detection is performed down to a 5σ threshold using the updated PSF models, followed by measurement of PSF and aperture fluxes. These catalogs represent the best single-epoch source characterization, but they are not intended for constructing light curves. For time-domain analysis, we recommend using the forced photometry tables described in §4.6.2

4.5. *Coaddition Processing*

4.5.1. *Coaddition*

Only exposures with a seeing better than 1.7 arcseconds FWHM are included in the deep coadded images. For the template coadds, only the top third of visits with the best seeing are used, resulting in an even tighter image quality cutoff for the template coadds.

Exposures with poor PSF model quality, identified using internal diagnostics, are excluded to prevent contamination of the coadds with unreliable PSF estimates. The remaining exposures are combined using an inverse-variance weighted mean stacking algorithm. To mitigate transient artifacts before coaddition, we apply the artifact rejection procedure described in AlSayyad (2019) that 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 common pixel grid ("warps") and leverages their temporal behavior to distinguish persistent astrophysical sources from transient artifacts.

Artifact rejection uses both direct and PSF-matched warps, homogenized to a standard PSF of 1.8 arcseconds FWHM, consistent with the 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 less than a small percentage of the total exposures, with the threshold varying based on the number of visits, N, as follows:

- $N = 1$ or 2: threshold $= 0$ (no clipping).
- $N = 3$ or 4: threshold $= 1$.
- $N = 5$: threshold $= 2$.
- $N > 5$: threshold $= 2 + 0.03N$.

Identified transient regions are masked before coaddition, improving image quality and reducing contamination in derived catalogs.

4.5.2. *Coadd Processing*

Coadd-processing consists of detection, deblending, and measurement on coadds to produce object tables (§3.2). For each coadd in all six bands, we fit a constant background and performed source detection at a 5σ detection threshold. Detections across bands are merged in a fixed priority order, *irzygu*, to form a union detection catalog, which serves as input to deblending.

Deblending is performed using the Scarlet Lite algorithm, which implements the same model as Scarlet (Melchior et al. 2018), but operates on a single pixel grid. This allows the use of analytic gradients, resulting in greater computational speed and memory efficiency.

Source measurement is then performed on the deblended footprints in each band. Measurements are conducted in three modes: independent per-band measurements, forced measurements in each band, and multiband measurements. Most measurement algorithms operate through a single-band plugin system, largely as originally described in Bosch et al. (2018). These plugins run on a deblended image, which is generated by using the Scarlet model as a template to reweight the original noisy coadded pixel values. This effectively preserves the original image in regions where objects are not blended, while dampening the noise elsewhere.

Measurement algorithm outputs include object fluxes, centroids, and higher-order moments thereof like sizes and shapes.

A reference band is then chosen for each object based on detection significance and measurement quality using the same priority order as detection merging (*irzygu*) and a second round of measurements is performed in forced mode using the shape and position from the reference band to ensure consistent colors (Bosch et al. 2018). A variety of flux measurements are included in the object tables, from aperture fluxes and forward modeling algorithms.

Composite model (CModel) magnitudes are used to calculate the extendedness parameter, which functions as a star-galaxy classifier. Gaussian-aperture-and-PSF (GAaP Kuijken 2008; Kannawadi 2022) fluxes are provided to ensure consistent galaxy colors across bands. Sersic model fits are run on all available bands simultaneously (MultiProFit Taranu 2025). The resulting Sersic (Sérsic 1963; Sersic 1968) model fluxes are provided as an alternative to CModel and are intended to represent total galaxy fluxes. Like CModel, the Sersic model is a Gaussian mixture approximation to a true Sersic profile, convolved with a Gaussian mixture approximation to the PSF. CModel measurements use a double "shapelet" (Refregier 2003) PSF with a single shared shape, while the Sersic fits use a double Gaussian with independent shape parameters for each component. Sersic model fits also include a free centroid, with all other structural parameters shared across all bands. That is, the intrinsic model has no color gradients, but the convolved model may have color gradients if the PSF parameters vary significantly between bands.

Further details on the performance of these algorithms can be found in 5.6.

4.6. *Variability Measurement*

4.6.1. *Difference Imaging Analysis*

Difference Image Analysis (DIA) used the decorrelated Alard & Lupton image differencing algorithm (Reiss & Lupton 2016). We detected both positive and negative `DIASource` at 5σ in the difference image. Sources with footprints containing both positive and negative peaks were fit with a dipole centroid code.

We filter a subset of `DIASources` that have pixel flags characteristic of artifacts, non-astrophysical trail lengths, and unphysically negative direct fluxes. We performed a simple spatial association of `DIASources` into `DIAObjects` with a one arcsecond matching radius.

To meet the latency requirements for Alert Production, we initially developed a relatively simple Machine Learning reliability model: a Convolutional Neural Network with three convolutional layers, and two fully connected layers. The convolutional layers have a 5×5 kernel size, with 16, 32, and 64 filters, respectively. A max-pooling layer of size 2 is applied at the end of each convolutional layer, followed by a dropout layer of 0.4 to reduce overfitting. The last fully connected layers have sizes of 32 and 1. The ReLU activation function is used for the convolutional layers and the first fully connected layer, while a sigmoid function is used for the output layer to provide a probabilistic interpretation. The cutouts are generated by extracting postage stamps of 51×51 pixels centered on the detected source. The input data of the model consists of the template, science, and difference image stacked to have a tensor of shape (3, 51, 51). The model is implemented using PyTorch (Ansel et al. 2024). The Binary Cross Entropy loss function was used, along with the Adaptive Moment Estimation (Adam) optimizer with a fixed learning rate of 1×10^{-4}, weight decay of 3.6×10^{-2}, and a batch size of 128. The final model uses the weights that achieved the best precision/purity for the test set. Training was done on the SLAC National Accelerator Laboratory () Shared Scientific Data Facility () with an NVIDIA L40S GPU model.

The model was initially trained using simulated data from the second Data Challenge (DC2; (LSST Dark Energy Science Collaboration (LSST DESC) et al. 2021)) plus randomly located injections of PSFs to increase the number of real sources, for a total of 89,066 real sources. The same number of bogus sources were selected at random from non-injected DIASources. Once the LSSTComCam data was available, the model was fine-tuned on a subset of the data containing 183,046 sources with PSF injections. On the LSSTComCam test set, the model achieved an accuracy of 98.06%, purity of 97.87%, and completeness of 98.27%.

4.6.2. *Lightcurves*

To produce light curves, we perform multi-epoch forced photometry on both the direct visit images and the difference images. For lightcurves we recommend the forced photometry on the difference images (`psDiffFlux` on the ForcedSource Table), as it isolates the variable component of the flux and avoids contamination from static sources. In contrast, forced photometry on direct images includes flux from nearby or blended static objects, and this contamination can vary with seeing. Centroids used in the multi-epoch forced photometry stage are taken either from object positions measured on the coadds or from the DIAObjects (the associated DIASources detected on difference images).

This stage takes the longest in terms of integrated Central Processing Unit (CPU)-hours.

4.6.3. *Solar System Processing*

Solar system processing in DP1 consists of two key components: the association of observations (sources) with known solar system objects, and the discovery of previously unknown objects by linking sets of *tracklets*[17].

To generate expected positions, ephemerides are computed for all objects found in the Minor Planet Center orbit catalog using the **Sorcha** survey simulation toolkit (Merritt et al., in press)[18]. To enable fast lookup of objects potentially present in an observed visit, we use the **mpsky** package (Juric 2025). In each image, the closest DiaSource within 1 arcsecond of a known solar system object's predicted position is associated to that object.

Solar system discovery uses the **heliolinx** package of asteroid identification and linking tools (Heinze et al. 2023). The suite consists of the following tasks:

- Tracklet creation with **make_tracklets**
- Multi-night tracklet linking with **heliolinc**
- Linkage post processing (orbit fitting, outlier rejection, and de-duplication) with **link_purify**

The inputs to the **heliolinx** suite included all sources detected in difference images produces by an early processing of the LSSTComCam commissioning data, including some that were later rejected as part of DP1 processing and hence are not part of this DP1 release.

About 10% of all commissioning visits targeted the near-ecliptic field Rubin_SV_38_7 designed to enable asteroid discovery. Rubin_SV_38_7 produced the vast majority of asteroid discoveries, as expected, but a few were found in off-ecliptic fields as well.

Tracklet creation with **make_tracklets** used an upper limit angular velocity of 1.5 deg/day, faster than any main belt asteroid and in the range of many Near-Earth Object () discoveries. To avoid excessive false tracklets from fields that were observed many times per night, the minimum tracklet length was set to three and the minimum on-sky motion for a valid tracklet was set to five arcseconds.

The heart of the discovery pipeline is the **heliolinc** task, which connects ("links") tracklets belonging to the same object over a series of nights. It employs the HelioLinC3D algorithm (Eggl et al. 2020; Heinze et al. 2022), a refinement of the original HelioLinC algorithm of Holman et al. (2018). The **heliolinc** run tested each tracklet with 324 different hypotheses spanning heliocentric distances from 1.5 to 9.8 AU and radial velocities spanning the full range of possible bound orbits (eccentricity 0.0 to nearly 1.0). This range of distance encompasses all main belt asteroids and Jupiter Trojans, as well as many comets and Mars-crossers and some NEOs. Smaller heliocentric distances were not attempted here because nearby objects move rapidly across the sky and hence were not likely to remain long enough in an LSSTComCam field to be discovered. A clustering radius was chosen corresponding to 1.33×10^{-3} AU at 1 deprecated acronym for astronomical unit; use astronomical unit (au) instead (au) from Earth. Linkages produced by **heliolinc** are then post-processed with **link_purify** into a final non-overlapping set of candidate discoveries, ranked from highest to lowest probability of being a real asteroid based on astrometric orbit-fit residuals and other considerations.

[17] A tracklet is defined as two or more observations taken in close succession in a single night.

[18] Available at https://github.com/dirac-institute/sorcha

5. PERFORMANCE CHARACTERIZATION AND KNOWN ISSUES

In this section, we provide an assessment of the DP1 data quality and known issues. A summary of the Rubin DP1 key numbers and data quality metrics and is found in PERF-SUMMARYTABLE

5.1. *Sensor Anomalies and ISR*

In addition to the known detector features identified before LSSTComCam commissioning, most of which are handled by the ISR processing (see §4.2.1), we discovered a number of new types of anomalies in the DP1 data. Since no corrections are currently available for these anomalies, they are masked and excluded from downstream data products.

5.1.1. *Vampire Pixels*

Vampire pixels are visible on the images as a bright defect surrounded by a region of depressed flux, as though the defect is stealing charge from its neighboring pixels; they have been termed "vampire" defects. From studies on evenly illuminated images, vampires appear to conserve charge. Unfortunately, there's no clean way to redistribute this stolen flux, and so we have identified as many of them as possible and created manual defect masks to exclude them from processing. We have found some similar features on the ITL detectors on LSSTCam, and will use the same approach to exclude them.

5.1.2. *Phosphorescence*

Some regions were seen to contain large numbers of bright defects. On closer study, it appears that on some detectors a layer of photoresist wax was incompletely removed from the detector surface during production. As this wax is now trapped below the surface coatings, there is no way to physically clean these surfaces. If this wax responded to all wavelengths equally, then it would likely result in quantum efficiency dips, which might be removable during flat correction. However, it appears that this wax is slightly phosphorescent, with a decay time on the order of minutes, resulting in the brightness of these sources being dependent on the illumination of prior exposures. The worst of these regions were excluded with manual masks, but we do not expect to need to do this for LSSTCam.

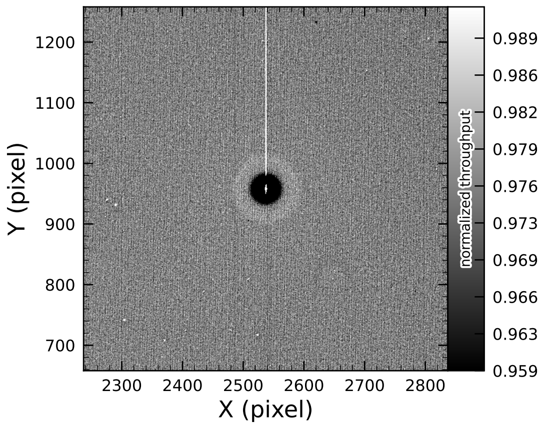

Figure 10. A large *vampire pixel* near the center of R22_-S11, as seen on the r-band flat.

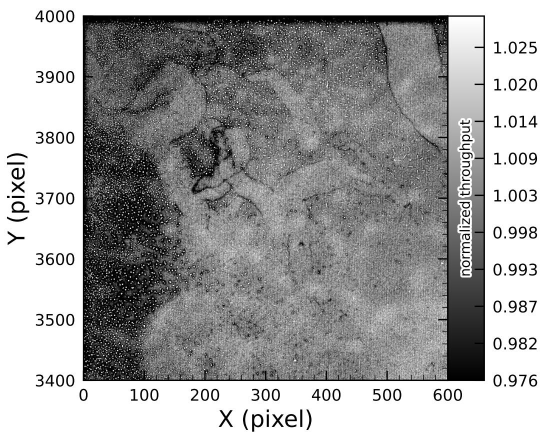

Figure 11. The top left corner of R22_S01 in the g-band flat, showing the many small defect features that are caused by the remnant photoresist wax. A single large defect box masks this region from further analysis to prevent these features from contaminating measurements.

5.1.3. *Crosstalk*

We use an average crosstalk correction based on laboratory measurements with LSSTCam. These average corrections performed better than expected, and so have been used as-is for DP1 processing. There are, however, some residual crosstalk features present post-correction, with a tendency towards over-subtraction.

5.1.4. *Bleed Trails*

Bleed trails from saturated sources were expected on LSSTComCam, but they appear in more dramatic forms than was expected. As a bleed trail nears the serial register, it fans out into a "trumpet" shaped feature. Although bright, these features do not have consistently saturated pixels, and were ignored in the first on-sky processing. We have

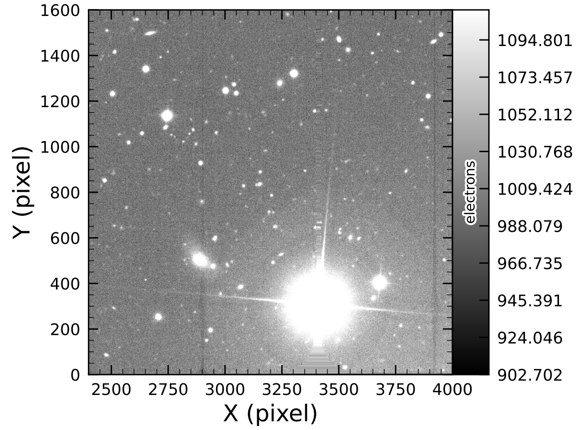

Figure 12. An example of a bright star with over-subtracted crosstalk residuals visible on neighboring amplifiers to both sides (exposure 2024120600239, detector R22_S02). The horizontal banding stretching from the center of the star shows the interpolation pattern covering the saturated core and the ITL edge bleed near the serial register.

since developed the means to programmatically identify and mask these features, which we have named "edge bleeds."

Saturated sources can create a second type of bleed, where the central bleed drops below the background level. The depressed columns along these trails extend across the entire height of the detector, crossing the detector mid-line. We developed a model for these to identify which sources are sufficiently saturated to result in such a trail, which is then masked. As these kind of trails appear only on the ITL detectors, we've named these features "ITL dips."

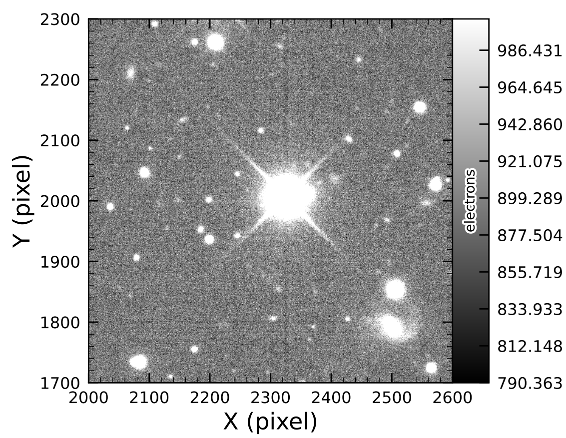

Figure 13. A bright star showing the "ITL dip" phenomenon, in which a dark trail extends out from the star to the top and bottom edges of the detector (exposure: 2024121000503, detector: R22_S21).

5.2. *PSF Models*

To characterize PSF performance, we use the second moments measured on PSF stars and on the PSF model via the Half-Second Moment (HSM) method (Hirata & Seljak 2003 and Mandelbaum et al. 2005), all expressed in the camera's pixel frame. Given the second-moment matrix elements I_{xx}, I_{yy}, and I_{xy}, we define:

$$T = I_{\mathrm{xx}} + I_{\mathrm{yy}}$$

$$e^1 = \frac{I_{\mathrm{xx}} - I_{\mathrm{yy}}}{T}$$

$$e^2 = \frac{2 I_{\mathrm{xy}}}{T}.$$

We denote T_{PSF}, e^1_{PSF}, and e^2_{PSF} for measurements on the PSF stars, and T_{model}, e^1_{model}, and e^2_{model} for the PSF model. Two variants are compared:

- Piff with second-order polynomial interpolation (default in science pipelines); and
- Piff with fourth-order polynomial interpolation (final DP1 PSF).

Table 4 summarizes each model's ability to reconstruct the mean T, e^1, and e^2 on LSSTComCam. Piff shows a negative residual bias in size. We will explore this further by plotting $\delta T/T$ versus magnitude (binned by color) in Fig. 16.

Another way to assess PSF performance is to examine the average across visits of $\delta T/T$ projected onto focal-plane coordinates (Figure 14). Piff shows strong spatial correlations, with a systematic offset that matches Table 4. It is the existence of these spatial structures that motivated raising the interpolation order to four, except in the u-band. Although not shown in Figure 14, third-order polynomial interpolation still exhibited residual structure. A fifth-order polynomial interpolation would require more stars than are available on some CCDs to adequately constrain the model while offering only marginal gains. Preliminary analysis of LSSTCam data in the laboratory at SLAC shows that the ITL sensors exhibit the same pattern. The sensor's $\delta T/T$ is fully correlated with the height variation across the LSSTCam ITL sensors, which explains this behavior. Future data processing will account for this height variation directly in the PSF model.

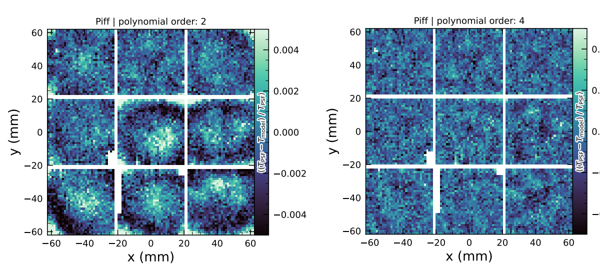

Figure 14. Average across all visits of $\delta T/T$ for different PSF modeling on LSSTComCam. Average is computed on a bin size of 120 pixels.

Another way to look at the PSF modeling quality is via whisker plots of the PSF second and fourth moments and their modeling residuals projected on a part of the sky. In addition to the second moment, the spin-2 fourth moments, $e^{(4)}$, is defined as:

$$e^{(4)}_1 = M_{40} - M_{04}$$

$$e^{(4)}_2 = 2\,(M_{31} - M_{13})\,,$$

where M_{pq} are the standardized higher moments as defined in Zhang et al. (2023) measured on stars and PSF models. Figure 15 shows the whisker plots of e, $e^{(4)}$ (top rows), and δe, $\delta e^{(4)}$ in the ECDFS field. The direction of the whiskers represents the orientation of the shape, while the length, modulated by the red bar, represents the amplitude $|e|$ or $|e^{(4)}|$. We observe coherent patterns in both the PSF moments and the residuals, the latter of which warrants further investigation if it persists in future data releases.

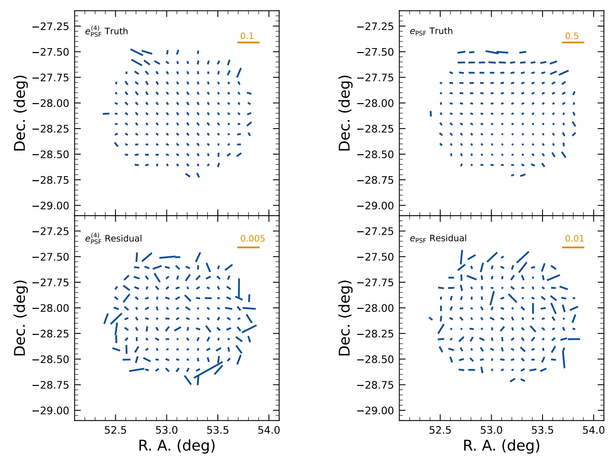

Figure 15. Whisker plot on ECDFS field for e, $e^{(4)}$ and δe, $\delta e^{(4)}$.

Another characterization of PSF-modeling performance is to look at $\delta T/T$ versus stellar magnitude to reveal any PSF size–flux dependencies (Figure 16). We also repeat this analysis in color bins to probe chromatic effects. Fainter stars show a larger negative bias in PSF size compared to brighter ones. Binning by color uncovers a clear color dependence, as seen in DES (e.g., Jarvis et al. 2021). DP1 does not include the color correction implemented in Schutt et al. (2025). Post-DP1 tests added a color correction similar to Schutt et al. (2025): it reduced the color-dependent scatter in PSF size but did not eliminate the negative bias for faint sources. The cause of this residual remains unknown and is consistent with what is shown in Table 4.

As mentioned in Developers (2025), there are two important Piff features that were not used during DP1. First, PSF color dependence was not yet implemented but will be added in the next release of the Rubin Science Pipelines. Second, although the current Rubin software allows Piff to operate in sky coordinates (including WCS transformations), it does not yet correct for sensor-induced astrometric distortions (e.g., tree rings). That capability is also planned for future data releases.

Table 4. Comparison of observed and model residuals, across all visits and filters.

Quantity	**Observed**	**Piff O2**	**Piff O4**
	$\times 10^{-3}$	$\times 10^{-4}$	$\times 10^{-4}$
$\langle T \rangle$ (pixel^2)	11.366 ± 0.003		
$\langle e^1 \rangle$	$(-6.07 \pm 0.05) \times 10^{-3}$		
$\langle e^2 \rangle$	$(-4.57 \pm 0.05) \times 10^{-3}$		
$\langle e \rangle$	$(8.794 \pm 0.004) \times 10^{-2}$		
$\langle \delta T/T \rangle$		-4.0 ± 0.2	-5.0 ± 0.2
$\langle \delta e^1 \rangle$		0.6 ± 0.1	0.5 ± 0.1
$\langle \delta e^2 \rangle$		0.0 ± 0.1	0.0 ± 0.1

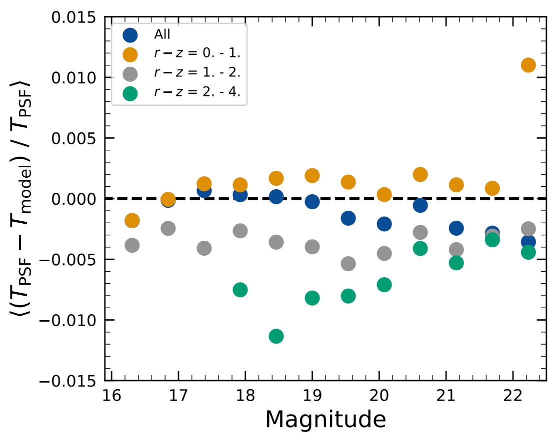

Figure 16. Binned $\delta T/T$ as a function of magnitude across all visits and filters and binned in different colors.

5.3. *Astrometry*

To characterize astrometric performance, we evaluate both internal consistency and agreement with an external reference. A primary measure of internal consistency is the repeatability of position measurements for the same object. We associate isolated point sources across visits and compute the Root-Mean-Square (RMS) of their fitted positions. Figure 17 shows the median per-tract astrometric error for all isolated point sources, both after the initial calibration and after the final calibration, which includes proper motion corrections. The results indicate that the astrometric solution is already very good after the initial calibration. Global calibration yields only modest improvement, likely due to the short time span of DP1 and the minimal distortions in the LSSTComCam. In the main survey, the longer time baseline and greater distortions near the LSSTCam field edges will make global calibration more impactful.

An additional metric of internal consistency is the repeatability of separations between objects at a given distance. To calculate this, we find pairs of objects at a given distance from each other, then calculate their separation in each visit in which they appear. The scatter in these distances then gives us a measure of the internal consistency of the astrometric model. The median value for each tract for objects separated by approximately 5 arcmin after the final calibration, i.e., AM1 from Ivezić & The LSST Science Collaboration (2018), is given in Figure 17. These values are already approaching the design requirement of 10 mas.

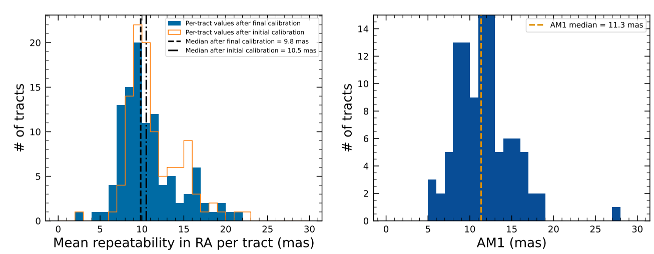

Figure 17. (a) Mean per-tract astrometric repeatability of measurements of isolated point sources in Rapid Analysis (RA) (b) Median per-tract repeatability in separations between isolated point sources 5 arcmin apart.

Finally, we consider the median separation between sources not included in the astrometric fit and associated objects from a reference catalog. For this, we use the Gaia DR3 catalog, with the object positions shifted to the observation epoch using the Gaia motion parameters. Figure 18 shows the median separation for each visit in the r-band in tract 4849.

The calculated values are almost all within 5 mas, well below the design requirement of 50 mas for the main survey.

By looking at the astrometric residuals, we can assess whether there are distortions not accounted for by the astrometric model. In some cases, the residuals in a single visit show behavior consistent with atmospheric turbulence, as shown in Figure 19. As in Léget et al. (2021) and Fortino et al. (2021), this is characterized by a curl-free gradient field in the two-point correlation function of the residuals (E-mode). However, as seen in Figure 20, the residuals in many visits also have correlation functions with a non-negligible divergence free B-mode, indicating that some of the remaining residuals are due to unmodeled instrumental effects, such as rotations between visits.

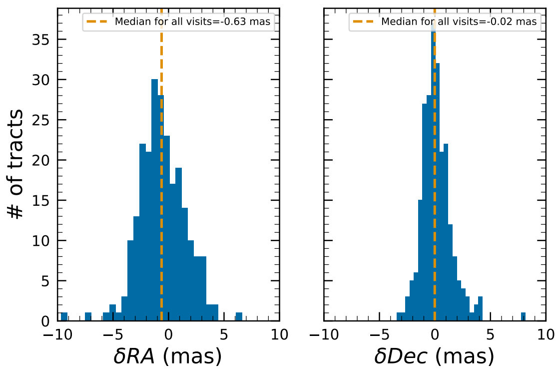

Figure 18. Median absolute offset for all visits in r-band in tract 4849. The offset is the difference between the position of isolated point sources that were reserved from the astrometric fit and matched objects from the Gaia DR3 catalog.

We can see unmodeled camera distortions by stacking the residuals over many visits as a function of the focal plane position. Figure 21 shows the median residuals in x and y directions for 1792 visits. Spatial structures are evident at the CCD level, along with the mid-line break in the y-direction residuals.

Further stacking all the detectors makes certain effects particularly clear. Figure 22 shows distortions very similar to those measured for an LSSTCam ITL sensor in a laboratory setting in Esteves et al. (2023).

5.4. *Photometry*

Eli: Photometry subsection is still needed

Repeatability on calibration star is .... and on psf flux stars is ....

5.5. *Detection Completeness on Coadds*

We characterize completeness by injecting synthetic sources into coadded images, and by comparing to external catalogs. In both cases, we use a greedy, probabilistic matching algorithm, whereby reference objects are matched in order of descending brightness to the most likely target within a $0.5''$ radius.

We inject sources in 12 of the patches of the ECDFS region with the deepest coverage. The input catalog contains stars and galaxies from part of the Data Challenge 2 (Dark Energy Science Collaboration ()) () simulations (LSST Dark Energy Science Collaboration (LSST DESC) et al. 2021), where the galaxies consist of an exponential disk and de Vaucouleurs (de Vaucouleurs 1948, 1953) bulge. To avoid deblender failures from excessive increases in object density, stars whose total flux (i.e., summed across all six bands) is brighter than 17.5 $\mathrm{mag_{AB}}$ are excluded, as are galaxies whose total flux is brighter than 15 $\mathrm{mag_{AB}}$ or fainter than 26.5 $\mathrm{mag_{AB}}$. Half of the remaining objects are selected for injection.

Figure 23 shows completeness as a function of magnitude for these injected objects. The completeness estimates are comparable to results from matching external catalogs. The Hubble Legacy Field catalog (Whitaker et al. 2019; Illingworth et al. 2016) reaches 50% completeness at 26.13 $\mathrm{mag_{F775W}}$, approximately 0.4 magnitudes fainter; this is roughly equivalent to 25.83 $\mathrm{mag_i}$ from differences in matched object magnitudes. Similarly, completeness drops below 90% at 23.80$\mathrm{mag_{VIS}}$ matching to Euclid Q1 (Euclid Collaboration et al. 2025) objects, equivalent to about 23.5 $\mathrm{mag_i}$. The Euclid imaging is of comparable (or shallower) depth, so magnitude limits at lower completeness percentages than 90% are unreliable, whereas the HST images cover too small (and irregular) of an area to accurately characterize 80-90% completeness limits.

At the 80% completeness limit, nearly 20% of objects, primarily injected galaxies, are incorrectly classified as stars based on the refExtendedness parameter, which indicates whether a source is more likely to be a point source or an extended source. Similarly, the fraction of correctly classified injected stars drops to about 50% at 23.8 $\mathrm{mag_i}$ (90% completeness).

There are several caveats for this analysis. The selection of objects for matching in any catalog is not trivial. Some fraction of the detections are either artifacts (particularly close to diffraction spikes around bright stars) or otherwise spurious. Additionally, some objects lie in masked regions of one survey but not another, which has not been accounted for. For injected source matching, the reference catalog does not include real on-sky objects. For this reason, we do not quote specific figures for purity; however, based on prior analyses of the DC2 simulations, purity is generally higher than completeness at any given magnitude.

5.6. *Flux Measurement*

Figure 24 shows i-band magnitude residuals for CModel and Sersic measurements using the matched injected galaxies described in 5.5. Similar behavior is seen in other bands. Sersic fluxes show reduced scatter and are more accurate on average for galaxies brighter than 22.5 $\mathrm{mag_i}$, though CModel's are less biased, median residuals are slightly closer to zero. For fainter objects, Sersic fluxes are more biased and less accurate. The magnitude of this bias is considerably larger than previously seen in simulated data and is being investigated. Aperture fluxes - including Kron and Gaussian Aperture and PSF () - are not shown as they are not corrected to yield total fluxes and thus are not recommended for use as total galaxy magnitudes.

Figure 24 shows $g-i$ color residuals versus r-band magnitude for the same sample of galaxies as Figure 24. For this and most other colors, GAaP (with a $1''$ aperture) and Sersic colors both yield lower scatter; however, the CModel colors have the smallest bias. Curiously, the GAaP bias appears to be magnitude-dependent, whereas the Sersic bias remains stable from $19 < r < 26$. Any of these color measurements are suitable for use for deriving quantities like photometric redshifts, stellar masses, etc.

In addition to photometry, some algorithms include measurements of structural parameters like size, ellipticity, and Sersic index. One particular known issue is that many (truly) faint objects have significantly overestimated sizes and fluxes, as was also seen in the Dark Energy Survey (Bechtol et al. 2025) and dubbed "super-spreaders". These super-spreaders

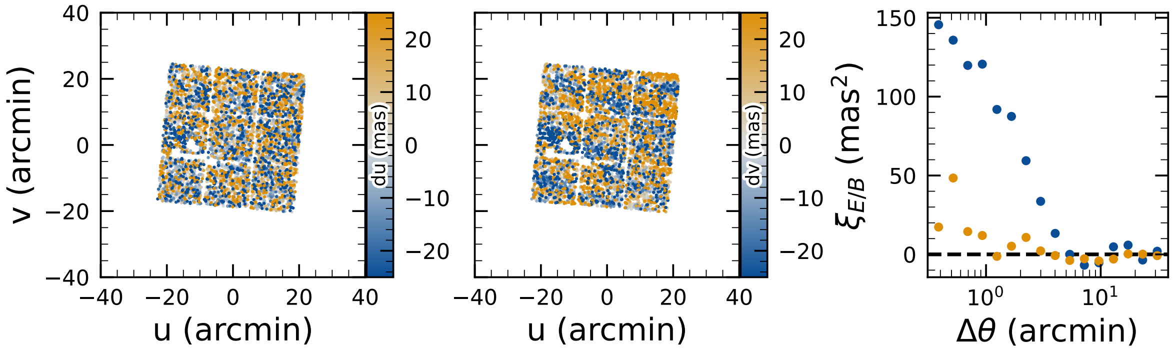

Figure 19. Residuals in *du* (left panel) and *dv* (center panel) directions, with the E and Byte (8 bit) (B)-modes of the two-point correlation function (right panel). The residuals show a wave-like pattern characteristic of atmospheric turbulence, and there is significant E-mode and negligible B-mode in the correlation function.

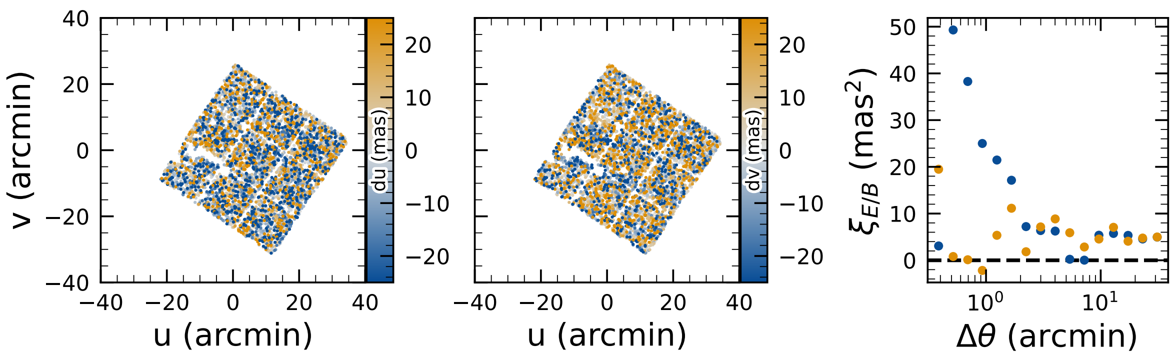

Figure 20. Residuals in *du* (left panel) and *dv* (center panel) directions, with the E and B-modes of the two-point correlation function (right panel). There are coherent residuals, but without the wave-like patter seen in Figure 19, and the correlation function has significant values for both E and B-modes.

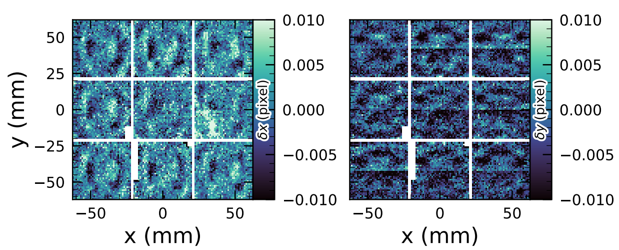

Figure 21. Median residuals as a function of focal plane position in *dx* (left panel) and *dy* (right panel) directions

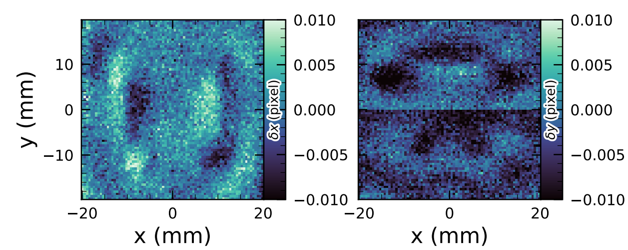

Figure 22. Median residuals as a function of pixel position in *dx* (left panel) and *dy* (right panel) directions

contribute significantly to overestimated fluxes at the faint end, and are particular problematic for the Kron algorithm (Kron 1980), which is not recommended for general use.

As mentioned in §4.5, the Sersic fits include a free centroid, which is initialized from the fiducial centroid of the object. Preliminary analyses of matched injected objects suggest that the galaxy astrometry residuals are somewhat smaller, and so users of the Sersic photometry should also use these centroid values (if needed). One caveat is that for faint objects and/or in crowded regions with unreliable deblending, free centroids can drift significantly and potentially towards other objects, so objects with large differences between the fiducial and Sersic astrometry should be used with caution.

5.7. *Differential Chromatic Refraction*

Differential Chromatic Refraction (DCR) occurs when light passes through Earth's atmosphere, refracting more

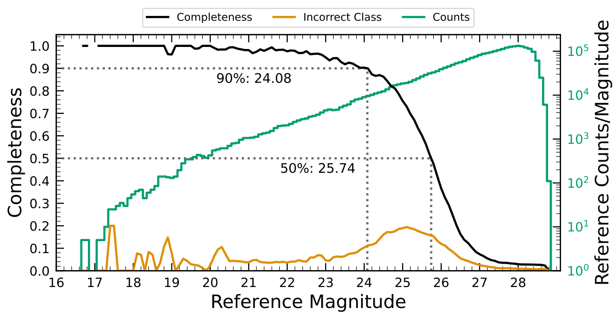

Figure 23. Completeness as a function of i-band CModel magnitude for DC2-based injections into a portion of the ECDFS field.

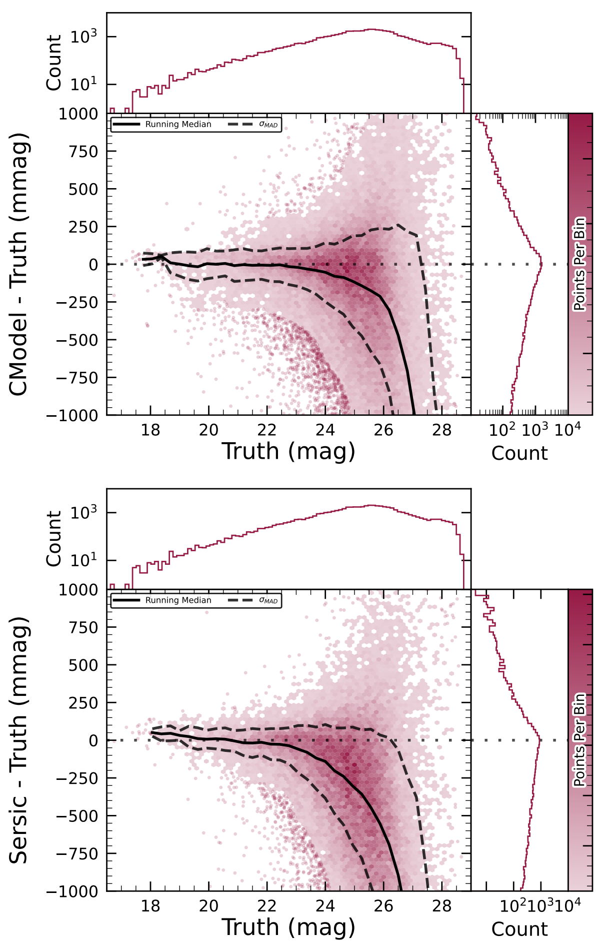

Figure 24. Magnitude residuals for matched injected galaxies with the CModel and Sersic algorithms.

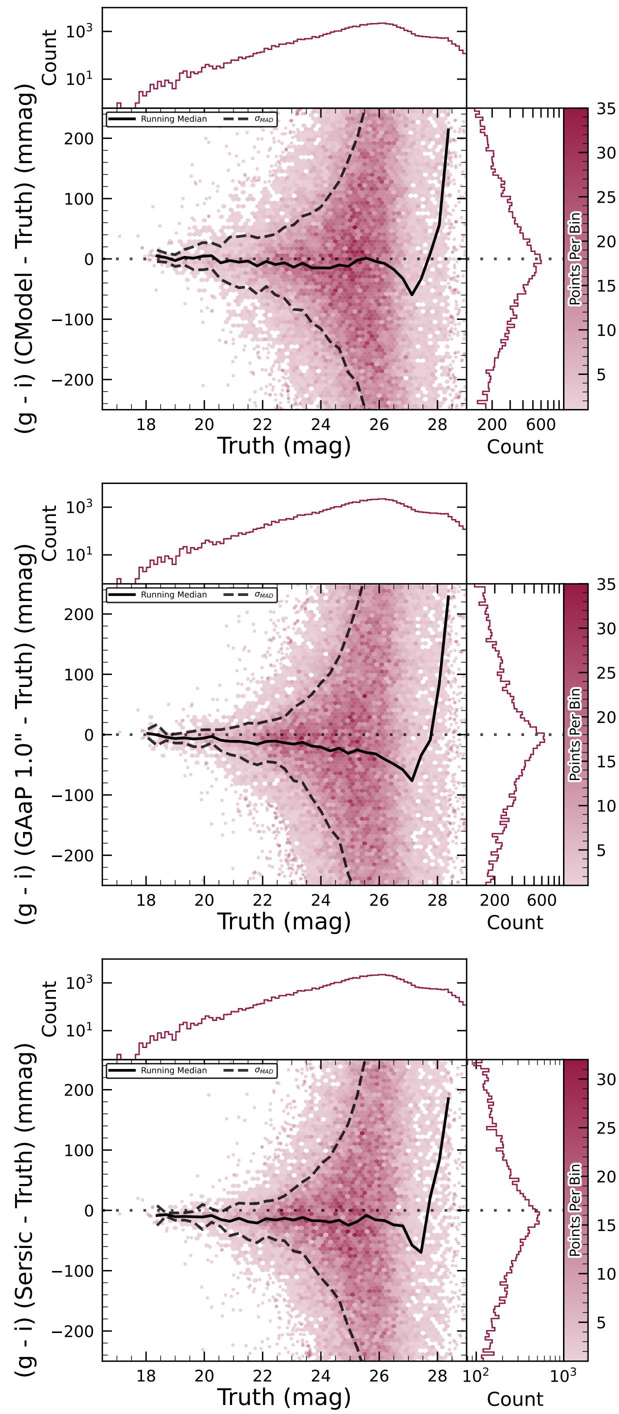

Figure 25. $g-i$ color residuals versus injected r-band magnitude for matched galaxies with the CModel, GAaP and Sersic algorithms.

for shorter wavelengths, which causes blue light to appear shifted closer to the zenith. This wavelength-dependent effect results in the smearing of point sources along the zenith direction, specifically parallel to the parallactic angle. The DCR effect is observable in LSSTComCam data, particularly in the angular offset versus g-i band magnitude difference plots Figure 26 which contains all direct sources with SNR > 10 from 41 visits from November 26, 2024. When looking at data perpendicular to the parallactic angle, sources show no DCR effect (as expected), forming a clear vertical distribution on the hexbin plots.

In contrast, sources aligned with the parallactic angle exhibit a tilted, linear distribution, clearly demonstrating the relationship between angular offset and the $g-i$ band magnitude difference, thereby providing a visual indication of the DCR effect.

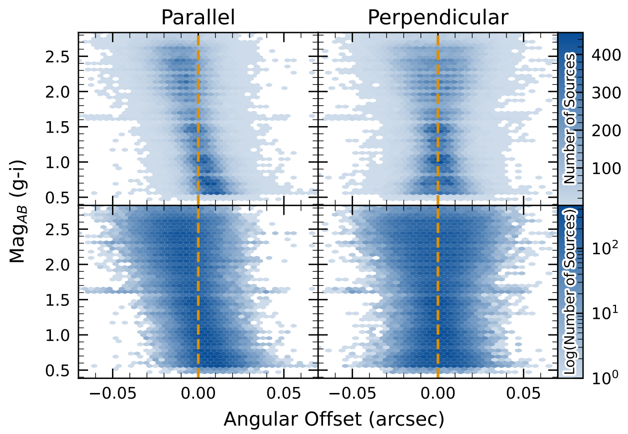

Figure 26. Visualization of Differential Chromatic Refraction (DCR) observed in the LSSTComCam commissioning campaign. The $g-i$ color is computed for every source in the reference catalog that is matched to a direct source in the science image, and the binned density for the full survey is plotted against the angular offset between the reference and detected positions. The angular offset is projected along coordinates parallel and perpendicular to the parallactic angle of the observation, and shows a characteristic correlation along the parallel axis with no correlation along the perpendicular axis. The orange vertical dashed line indicates the expected $g-i$ magnitude distribution at zero angular offset, while the green 'x' marks the average $g-i$ magnitude of the plotted sources.

5.8. *Difference Imaging Purity*

We assessed the performance of image differencing using human vetting and source injection (§5.9.3). Members of the DP1 team labeled more than 9500 DIASource image triplets consisting of cutouts from the science, template, and difference images. We classified these into various real and artifact categories. The raw real:bogus ratio was roughly 9:1. Bright stars are the main source of artifacts. Correlated noise, primarily in u and g bands, also leads to spurious detections near the threshold. We expect to be able to mitigate these effects for LSSTCam.

Applying a reliability threshold improves the purity of transients but not variable stars; technical limitations at the time of model training prevented injection of variable stars into the synthetic training set. Reliability models for LSSTCam data will be trained on a wider range of input data.

5.9. *Solar System*

5.9.1. *Asteroid Linking Performance*

DP1 performance evaluation of asteroid linking focused on demonstrating discovery capability. The solar system discovery pipeline produced 269,581 tracklets, 5,691 linkages, and 281 post-processed candidates.

We performed a conservative manual investigation of these 281 candidates, producing a curated list of 93 probable new asteroid discoveries. As described in Section 4.6.3, post processing of the `heliolinc` output with `link_purify` produced a final set of 281 candidate linkages, ranked with the most promising candidates first. Using `find_orb` (Gray 2025), we derived orbit fits for each candidate, sorting the resulting list by χ^2_{dof}, the quality of the fit. Manual inspection of the linkages indicated that those ranked 0–137 corresponded to unique real asteroids; ranks 138–200 contained additional real objects intermixed with some spurious linkages; and ranks higher than 200 were essentially all spurious. This analysis indicates that it will be possible to identify cuts on quality metrics like χ^2 to derive discovery candidate samples with high purity; determining the exact quantitative cut values require more data with LSSTCam. We next removed all observations matched to known asteroids (using Minor Planet Center ()'s MPChecker service), reducing the number of candidates to 97. Of these, four had strong astrometric and/or photometric outliers, likely due to self-subtraction in difference images due to the unavoidable limitations of template generation from the limited quantity of data available from LSSTComCam. We suspect these four linkages do correspond to real objects, but have chosen to discard them out of an abundance of caution. The remaining 93 were submitted to the Minor Planet Center and accepted as new discoveries, demonstrating the LSST pipelines are able to successfully discover new solar system objects.

Jake: We should cite the MPEC with discoveries, once we do submit and the MPEC becomes available

5.9.2. *Asteroid Association Performance*

Solar system association associated 5988 DiaSources to 431 unique solar system objects.

Jake: Update this after table update!

These include 3,934 DiaSources to 338 already-known MPC objects and 2,054 DiaSources to the 93 discoveries newly-discovered objects. Association also picked up an additional 143 detections of newly discovered objects.

Jake: This too - new parameter in notebook.

These were not originally found by the discovery pipelines as they didn't satisfy the number and/or maximum time span requirements to form tracklets.

The astrometric residuals of known asteroid association are shown in Figure 27.

Jake: Todo:

Astrometric precision for solar system sources is excellent, the majority of objects detected within $0''.1$ of their expected positions. Taking the unsigned median residuals to search for biases, we find that previously-known objects have mean

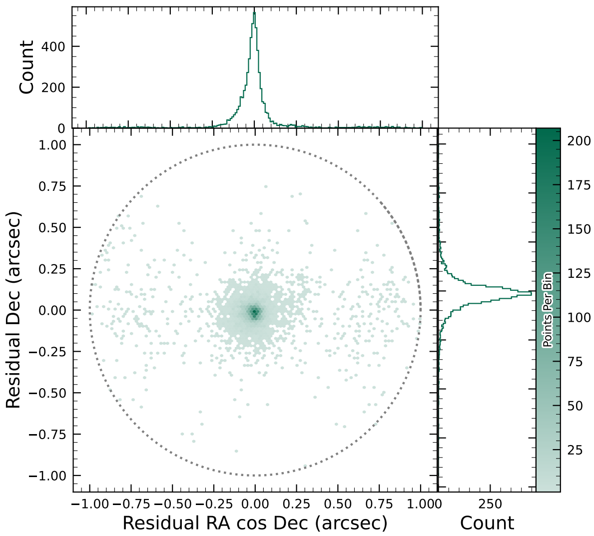

Figure 27. Astrometric residuals between expected and observed positions of SSOs in DP1. The median residuals are 0.″001 and −0.″016 in R.A./Dec direction, with the standard deviations of 0.″19 and 0.″10, respectively. No detectable detectable systematic offset from zero indicates there are no major errors in either timing or astrometry delivered by the Rubin system. The wider scatter in the RA-direction is due to objects whose measured orbital elements are less well constrained, translating to larger along-track positional errors in the predicted positions.

residuals of 0.″001 and −0.″016 in the RA and Dec directions respectively, while newly-discovered objects have mean residuals of −0.″035 and −0.″010 in the RA and Dec directions, respectively. These mean residuals are small enough to eliminate the possibility of a timing offset greater than the second-scale shutter motion (which is uncharacterized for LSSTComCam).

5.9.3. *Detection Completeness on Difference Images*

We assess the performance of our difference imaging pipeline using synthetic source injection on the science images prior to differencing. We construct a catalog of injected sources by joining two different samples of point sources, a set of hosted sources to emulate transients in galaxies and second set of hostless.

The hosts are selected from the pipeline source catalog that is produced upstream by imposing a cut in their extendedness measurement, and selecting $N_{\rm src} = \min(100, N \times 0.05)$ of the available sources per detector. For each host we pick a random position angle and radius using its light profile shape, and also a random value of brightness for the injected source, with magnitudes higher than the host source. The hostless sources instead have random positions in the CCD focal plane, and with magnitudes chosen from a random uniform distribution with $20 \geq m \geq m_{lim} + 1$ with m_{lim} the limiting magnitude of the image.

We used the LSST package `source_injection` to include these sources into our test images, we performed a coordinate cross-match task, with a threshold of 0.″5 to find which of these sources were detected and which were lost, enabling the calculation of a set of performance metrics.

In Figure 28 we show the detection completeness as function of the SNR, for sources in the ECDFS field, for filters *griz*. We observe a completeness $> 95\%$ for sources with SNR> 6, with mean completeness $\simeq 99\%$ and standard deviation of $\simeq 0.7\%$. In Figure 29 we show the distribution of

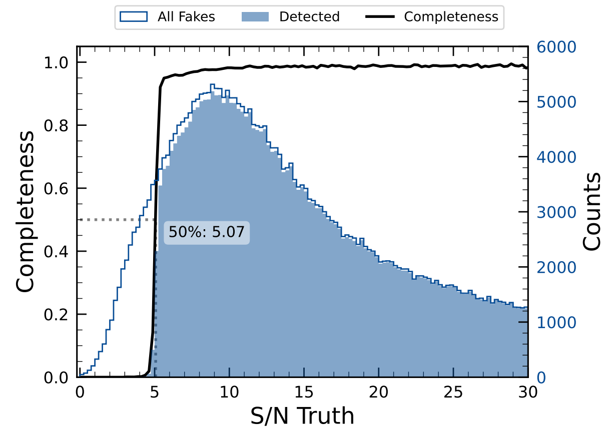

Figure 28. The difference image detection completeness for injected sources in the ECDFS field, for filters *griz*, as function of the estimated signal to noise ratio S/N. This completeness is the ratio between the found fake sources (shaded histogram) and all the sources (solid line). The horizontal dashed line represents where the 50% completeness level is reached, at approximately S/N $\simeq$ 5.07.

the residuals of the recovered sky coordinates for the detected synthetic sources. The marginal distributions are both centered at zero, and they are compatible with normal distributions $\mathcal{N}(0, 0''.04)$. In Figure 30 we show the recovered magnitudes for our detected synthetic sources in the i filter, using PSF photometry on the difference images, and also show marginal distributions of the true magnitudes for fake sources, and the residuals on the left, split into hosted and hostless. Our flux measurements are accurate within a wide range of magnitudes, for both hosted and hostless synthetic sources. We obtain that for true $m_i < 22.2$, the median PSF magnitudes residuals are < 0.1. When considering the flux pulls $\delta = (f - f_{\rm True})/\sigma_f$ for PSF flux f and error σ_f, we find that $|\langle\delta\rangle| < 0.1$, and $\sigma_\delta < 1.1$ for $m_i < 21.6$.

5.10. *Crowded Fields*

Two of the seven DP1 target fields exhibit high stellar density, 47 Tucanae and the Fornax dwarf galaxy. 47 Tucanae was chosen as an initial stress test for the science pipelines processing. The Fornax dwarf galaxy also exhibits high stellar density, particularly in its central regions.

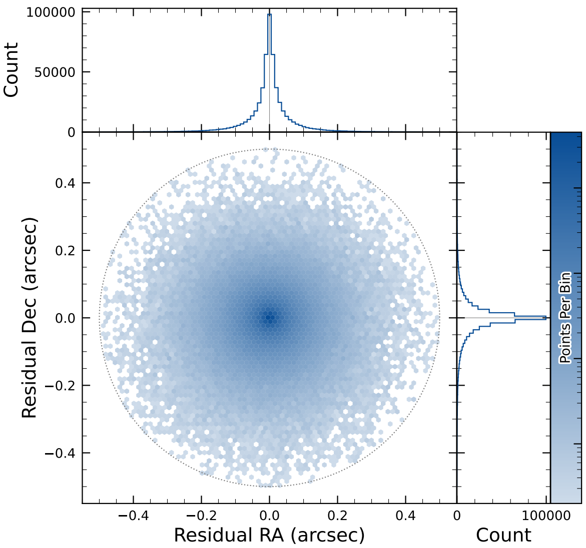

Figure 29. Coordinate residuals for detected synthetic sources in difference images, between recovered and true position of the sources in the ECDFS field. In the top and right panels we include the histogram of these offsets. The circle reflects the matching radius of 0″.5.

Yusra: Explain where the pipelines broke down. and how the performance is different in the 2 crowded fields

6. RUBIN SCIENCE PLATFORM

The RSP (Jurić et al. 2019) is a powerful, cloud-based environment for scientific research and analysis of petascale-scale astronomical survey data. It serves as the primary interface for scientists to access, visualize, and conduct next-to-the-data analysis of Rubin and LSST data. The RSP is designed around a "bring the compute to the data" principle, eliminating the need for users to download massive datasets. Although DP1 is comparable in size (3.5 TB) to existing survey datasets, future LSST datasets will be larger and more complex, making it crucial to co-locate data and analysis for effective scientific discovery.

The RSP provides users with access to data and services through three distinct user-facing Aspects: a *Portal*, which facilitates interactive exploration of the data; a JupyterLab-based *Notebook* environment for data analysis using Python; and an extensive set of *Application Programming Interfaces (APIs)* 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 data products described in §3 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 31 shows the Rubin Science Platform landing page in the Google cloud.

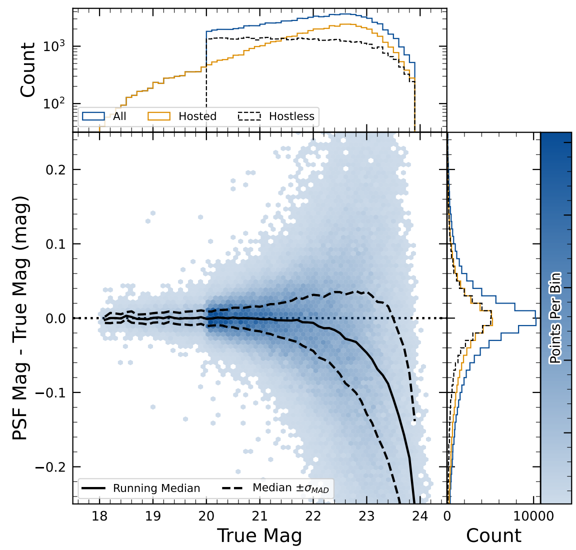

Figure 30. Magnitude residuals for PSF photometry on difference images for ECDFS field in i for detected fake sources. In black solid and dashed lines: the running median, and the mean absolute deviation. Top panel: the distribution of true magnitudes for hostless and hosted fakes sources. Right panel: the distribution of magnitude residuals for hostless and hosted sources.

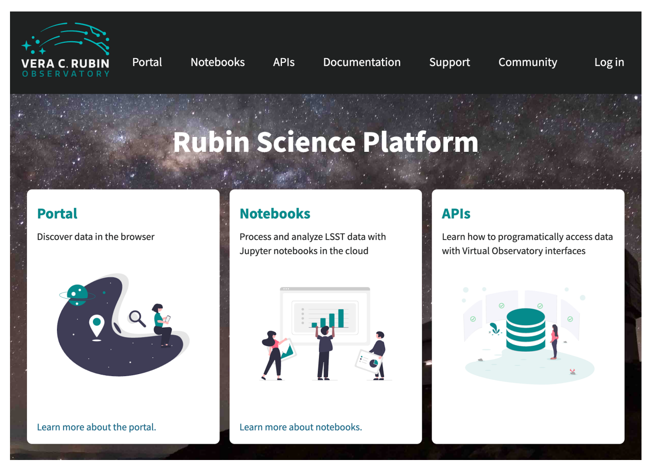

Figure 31. The Rubin Science Platform landing page showing the thress Aspects as well as links to documentation and support information.

The RSP is supported by a number of back-end services, including databases, files, and batch computing. Support for collaborative work through shared workspaces is also included in the RSP.

A preview of the RSP was launched on Google Cloud in 2022, operating under a shared-risk model to support Data Preview 0 (O'Mullane et al. 2024a). This allowed the community to test the platform, begin preparations for science, and provide valuable feedback to inform ongoing development. It was the first time an astronomical research environment was hosted in a cloud environment. The DP1 release brings major updates to RSP services, enhancing scientific analysis capabilities. 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 DP1 release and provides an overview of planned future capabilities.

6.1. *Rubin Data Access Center*

The Rubin USDAC utilizes a novel hybrid on-premises-cloud architecture, which combines on-premises infrastructure at the USDF at SLAC with flexible and scalable resources in the Google cloud. This architecture has been deployed and tested using the larger simulated data set of DP0.2 (O'Mullane et al. 2024b).

In this hybrid model, 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. The majority of the static data products described in §3 are stored on-premises at the USDF to benefit from cost-effective mass storage and close integration with Rubin data processing infrastructure, also located at the USDF. For imaging data, the Data Butler (§6.2.2) provides the interface between the cloud-based users and data services, and the on-premises data. For catalog data, a cloud-based TAP client (§6.2.1) submits queries to the on-premises Qserv database cluster (§6.5) and retrieves the results. In the initial DP1 deployment, catalog data is hosted at the USDF while image data is stored in the cloud. The full hybrid model will be rolled out and further tested following the release of DP1.

The RSP features a single-sign-on authentication and authorization system to provide secure access for Rubin data rights holders (Blum & the Rubin Operations Team 2020)

6.2. *API Aspect*

The 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 within the LSST Science Pipelines.

At the time of the DP1 release, some IVOA services are unavailable, and certain data products are only accessible via the Butler. This section provides an overview of the available IVOA services and Butler access.

6.2.1. *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 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) (IVOA) query on the Object catalog via TAP yields identical results whether run from the Portal, Notebook, or an external client.

The following IVOA services are available at the time of the DP1 release:

- **Table Access Protocol (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 DP1 runs on the Rubin-developed Qserv database (§ 6.5), which hosts the core science tables described in §3.2, as well as the Visit database. It also provides image metadata in the IVOA ObsCore format via the standard `ivoa.ObsCore` table, making it an "ObsTAP" service (ObsTAP; Louys et al. 2017). The TAP service is based on the Canadian Astronomy Data Centre (CADC)'s open-source Java TAP implementation[19], modified for the exact query language accepted by Qserv. It currently supports a large subset of ADQL, with limitations documented in the data release materials (see §7.1) 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 SIAv2 (Simple Image Access Protocol, version 2; Jenness et al. 2024; Dowler et al. 2015) for discovering and accessing astronomical images based on metadata. For example, querying for all images in a given band over a particular sky region observed during a given period. SIAv2 is a REpresentational State Transfer (REST)-based protocol that supports the discovery and retrieval of image data. 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 2025), allowing the dataset to be retrieved or a cutout requested via Server-side Operations for Data Access (IVOA standard) ().

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

- **Image Cutout Service**: The Rubin Cutout Service (Allbery 2023, 2024) is based on the IVOA SODA (Server-side Operations for Data Access; 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. For DP1, The cutout service is a single cutout service only where N cutout requests will require N independent synchronous calls. We expect some form of bulk cutout service by mid 2026, approximately contemporaneously with DP2
- **HiPS Data Service**: An authenticated HiPS (Fernique et al. 2017) data service for seamless pan-and-zoom access to large-scale co-adds. It supports fast interactive progressive image exploration at a range of resolutions.
- **WebDAV**: A Web Distributed Authoring and Versioning (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.

6.2.2. *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 S3 object store; and
- the *Registry*: An Structured Query Language (SQL)-compatible database that stores metadata about the datasets in the data store, see §6.5.2.

For DP1, the Butler repository is hosted in the Google Cloud, using an (Amazon) Simple Storage Service (S3)-compatible store for datasets and a PostgreSQL database for the registry.

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 §3. 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 example, a dataset that represents a single raw image with detector 8 during the on-sky campaign on the night starting 2024-11-11 in the i band with exposure ID 2024111100074 would be represented as `dataId='exposure':2024111100074, 'band':'i', 'instrument':'LSSTComCam'` and is associated with the `raw` DatasetType. For a deep coadd on a patch of sky in the Seagull field, there would be no exposure dimensions and would instead the tract, patch and band would be specified as `dataId='tract':7850, 'patch': 6, 'band':'g', 'instrument':'LSSTComCam', skymap='lsst_cells_v1'` and is associated with the `deep_coadd` DatasetType.

The data coordinate is used to locate a dataset in multi-dimensional space, where dimensions are defined in terms of scientifically meaningful concepts, such as instrument, visit, detector or band. For example, a calibrated single-visit image (§3.1) has dimensions including band, instrument, and detector. In contrast, the visit table (§3.2), a catalog of all calibrated single-epoch visits in DP1, has only the instrument dimension. The main dimensions used in DP1 are listed, together with a brief description, in Table 5. To determine which dimensions are relevant for a specific dataset, the Butler defines dataset types, which associate each dataset with its specific set of relevant dimensions, as well as the associated Python type representing the dataset. The dataset type defines the kind of data a dataset represents. For example, a raw image (`raw`), a processed catalog (`object_forced_source`), or a sky map (`skyMap`).

Table 6 lists all the dataset types available via the Butler in DP1, together with the dimensions needed to uniquely identify a specific dataset and the number of unique datasets of each type. It is important to highlight a key difference between accessing catalog data via the TAP service versus the Butler. While the TAP service contains entire catalogs, many of the same catalogs in the Butler are split into multiple separate catalogs. This is partly due to how these catalogs are generated, but also because of the way data is stored within and retrieved from the Butler repository – it is inefficient to retrieve the entire `Source` catalog, for example, from the file system. Instead, because the `Source` catalog contains data for sources detected in the `visit_images`, there is one `Source` catalog in the Butler for each `visit_image`. Similarly, there is one `Object` catalog for each `deep_coadd`. All the catalogs described in §3.2, aside from the `CcdVisit`, `SSObject`, `SSSource`, and `Calibration` catalogs, are split within the Butler.

A dataset is associated with one or more *Run Collections*; logical groupings of datasets within the Butler system that were created or processed together by the same batch operation. Collections allow multiple datasets with the same data coordinate to coexist without conflict. Run Collections support flexible, parallel processing by enabling repeated analyses of the same input data using different configurations.

For DP1, a subset of the consolidated database contents (§6.5.3) is accessible through the Data Butler. However, not all metadata from the `Visit` table (§3.4) is available. The DP1 Butler is read-only; a writeable Butler is expected by mid-2026, around the time of DP2.

6.2.3. *Remote Programmatic Access*

The Rubin RSP API can be accessed from a local system by data rights holders outside of the RSP, by creating a user security token. This token can then be used as a bearer token for API calls to the RSP TAP service. This capability is especially useful for remote data analysis using tools such as Tool for OPerations on Catalogues And Tables (TOPCAT), as well as enabling third-party systems (e.g., Community Alert Brokers) to access Rubin data. Additionally, it sup-

Table 5. Descriptions of and valid values for the key data dimensions in DP1. YYYYMMDD signifies date and # signifies a single 0-9 digit.

Dimension	Format/Valid values	Description
day_obs	YYYYMMDD	A day and night of observations that rolls over during daylight hours.
visit	YYYYMMDD#####	A sequence of observations processed together; synonymous with “exposure” in DP1.
exposure	YYYYMMDD#####	A single exposure of all nine ComCam detectors.
instrument	LSSTComCam	The instrument name.
detector	0 - 8	A ComCam detector.
skymap	lsst_cells_v1	A set of tracts and patches that subdivide the sky into rectangular regions with simple projections and intentional overlaps.
tract	See Table 2	A large rectangular region of the sky.
patch	0 - 99	A rectangular region within a tract.
band	u, g, r, i, z, y	An astronomical filter.

Table 6. The name and number of each type of data product in the Butler and the dimensions required to identify a specific dataset.

Data Product	Name in Butler	Required Dimensions	Number in DP1
raw	raw	instrument, detector, exposure	16125
visit_image	visit_image	instrument, detector, visit	15972
deep_coadd	deep_coadd	band, skymap, tract, patch	2644
template_coadd	template_coadd	band, skymap, tract, patch	2730
difference_image	difference_image	instrument, detector, visit	15972
Source	source	instrument, visit	1786
Object	object	skymap, tract	29
ForcedSource	object_forced_source	skymap, tract, patch	636
DiaSource	dia_source	skymap, tract	25
DiaObject	dia_object	skymap, tract	25
ForcedSourceOnDiaObject	dia_object_forced_source	skymap, tract, patch	597
CCDVisit	visit_detector_table	instrument	1
SSObject	ss_object	–	1
SSSource	ss_source	–	1
Visit	visit_table	instrument	1
x			

ports remote development with local IDEs, allowing for more flexible workflows and integration with external systems.

6.3. *Portal Aspect*

The Portal Aspect provides an interactive environment for exploratory data discovery, query, filtering, and visualization of both image and catalog data, without requiring programming experience.

It enables users to search, visualize, and interact with large datasets through tools for catalog queries, image browsing, time series inspection, and cross-matching. The Portal is designed to support both exploratory data access and detailed scientific investigation.

The Portal is built on Firefly (Wu et al. 2019), a powerful web application framework developed by IPAC (Infrared Processing and Analysis Center). Firefly provides interactive capabilities such as customizable table views, image overlays, multi-panel visualizations, and linked displays between catalogs and images. Through Firefly, the Portal delivers a responsive and intuitive user experience, allowing users to analyze data visually while maintaining access to underlying metadata and query controls.

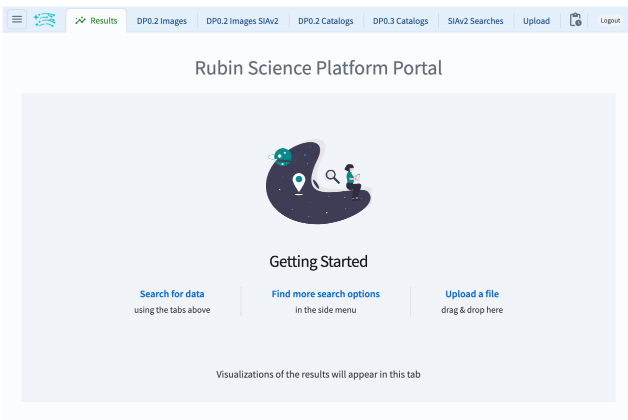

Figure 32. The Rubin Science Platform Portal Aspect, showing a DP1 image> Gregory: Replace with a real DP1 image

6.4. *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 (§4.1). The environment comes pre-installed with the pipelines and a broad set of widely used astronomical software tools, supporting immediate and flexible data analysis.

6.5. *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 DP1 and other Rubin databases is available online at https://sdm-schemas.lsst.io.

6.5.1. *Qserv*

The final 10-year 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. 2023) – 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: https://smonkewitz.github.io/scisql/) simplifying complex queries involving spherical geometry, statistics, and photometry. Qserv is built on robust production-quality components, including MariaDB (https://www.mariadb.org/) and XRootD (https://xrootd.org/). Qserv runs at the USDF and user access to catalog data is via the TAP service (§6.2.1). This enables catalog-based analysis through both the RSP Portal and Notebook Aspects.

Although the small DP1 dataset does not require Qserv's full capabilities, we nevertheless chose to use it for DP1 to accurately reflect the future data access environment and to gain experience with scientifically-motivated queries ahead of full-scale deployment. Qserv is open-source and available on GitHub: https://github.com/lsst/qserv.

6.5.2. *Butler Registry*

The Butler registry is a relational database that manages metadata and relationships between the various datasets in a data preview or release. For DP1, the registry is a PostgreSQL database.

6.5.3. *Consolidated Database*

The Consolidated Database (ConsDB) (Lim 2025) is an SQL-compatible database designed to store and manage metadata for Rubin Observatory science and calibration images. Metadata is recorded on a per-exposure basis and includes information such as the target name, pointing coordinates, observation time, physical filter and band, exposure duration, and environmental conditions (e.g., temperature, humidity, and wind speed). This key image metadata is also stored in the Butler Registry (§6.2.2), however the ConsDB stores additional information including derived metrics from image processing and information from the Engineering and Facility Database (EFD) transformed from the time dimension to the exposure dimension.

The ConsDB schema is organized into instrument-specific tables, e.g., LSSTComCam and LSSTCam, facilitating instrument-specific queries. Within the LSSTComCam schema, data is further structured into tables for individual exposures and detectors. An example query on the DP1 dataset might retrieve all visits within a specified time range in the r-band for a given DP1 target.

The ConsDB is hosted at the USDF. Following the initial release of DP1, a release of the DP1 exposure-specific ConsDB data will be made available through the RSP, and accessible externally via TAP. The detailed LSSTComCam schema can be found at: https://sdm-schemas.lsst.io/cdb_lsstcomcam.html

7. SUPPORT FOR COMMUNITY SCIENCE

The Rubin Observatory has a science community that encompasses thousands of individuals worldwide, with a broad range of experience and expertise in astronomy in general, and in the analysis of optical imaging data specifically.

Rubin's model to support this diverse community to access and analyze DP1 emphasizes self-help via documentation and tutorials, and employs an open platform for asynchronous issue reporting that enables crowd-sourced solutions. These two aspects of community support are augmented by virtual engagement activities. In addition, Rubin supports its Users Committee to advocate on behalf of the science community, and supports the eight LSST Science Collaborations.

All of the resources for scientists that are discussed in this section are discoverable by browsing the For Scientists pages of the Rubin Observatory website[20].

7.1. *Documentation*

The data release documentation for DP1 can be found at dp1.lsst.io. The contents include an overview of the LSSTComCam observations, descriptions of the data products (images and catalogs), and a high-level summary of the processing pipelines. Similar to the contents of this paper, but presented in a browsable, searchable webpage built with Sphinx[21], and written with a focus on applications of the data products to scientific analysis.

7.2. *Tutorials*

A suite of tutorials that demonstrate how to access and analyze DP1 using the RSP accompany the data release. Jupyter Notebook tutorials are available via the "Tutorials" drop-down menu within the Notebook aspect of the RSP. Tutorials for the Portal and API aspects of the RSP can be found in the data release documentation.

These tutorials are designed to be inclusive, accessible, clear, focused, and consistent. Their format and contents follow a set of guidelines[22] that are informed by industry standards in technical writing.

7.3. *Community Forum*

The venue for all user support is the Rubin Community Forum[23].

Questions about any and all aspects of the Rubin data products, pipelines, and services should be posted as new topics in the Support category. This includes beginner-level and "naive" questions, advanced scientific analysis questions, technical bug reports, account and data access issues, and everything in between. The Support category of the Forum is monitored by Rubin staff, who aim to respond to all new unsolved topics within 24 hours.

The Rubin Community Forum is built on the open-source Discourse platform. It was chosen because, for a worldwide community of ten thousand Rubin users, a traditional (i.e., closed) help desk represents a risk to Rubin science (e.g., many users with the same question having to wait for responses). The open nature of the Forum enables self-help by letting users search for similar issues, and enables crowd-sourced problem solving (and avoids knowledge bottlenecks) by letting users help users.

[20] https://rubinobservatory.org/

[21] https://www.sphinx-doc.org/

[22] Rubin's Guidelines for User Tutorials, https://rtn-045.lsst.io/.

[23] https://community.lsst.org/

7.4. *Engagement Activities*

A variety of live virtual and in-person workshops and seminars offer learning opportunities to scientists and students working with DP1.

- Rubin Science Assemblies (weekly, virtual, 1 hour): alternates between hands-on tutorials based on the most recent data release and open drop-in "office hours" with Rubin staff.
- Rubin Data Academy (annual, virtual, 3-4 days): an intense set of hands-on tutorials based on the most recent data release, along with co-working and networking sessions.
- Rubin Community Workshop (annual, virtual, 5 days), a science-focused conference of contributed posters, talks, and sessions led by members of the Rubin science community and Rubin staff

For schedules and connection information, visit the For Scientists pages of the Rubin Observatory website. Requests for custom tutorials and presentations for research groups are also accommodated.

7.5. *Users Committee*

This committee is charged with soliciting feedback from the science community, advocating on their behalf, and recommending science-driven improvements to the LSST data products and the Rubin Science Platform tools and services. Community members are encouraged to attend their virtual meetings and raise issues to their attention, so they can be included in the committee's twice-yearly reports to the Rubin Observatory Director.

The community's response to DP1 will be especially valuable input to DP2 and Data Release 1 (), and the Users Committee encourages all users to interact with them. For a list of members and contact information, visit the For Scientists pages of the Rubin Observatory website.

7.6. *Science Collaborations*

The eight LSST Science Collaborations are independent, worldwide communities of scientists, self-organized into collaborations based on their research interests and expertise. Members work together to apply for funding, build software infrastructure and analysis algorithms, and incorporate external data sets into their LSST-based research.

The Science Collaborations also provide valuable advice to Rubin Observatory on the operational strategies and data products to accomplish specific science goals, and Rubin Observatory supports the collaborations via staff liaisons and regular virtual meetings with Rubin operations leadership.

8. SUMMARY AND FUTURE RELEASES

Rubin Data Preview 1 (DP1) offers an initial look at the first on-sky data products and access services from the Vera C. Rubin Observatory. DP1 forms part of Rubin's Early Science Program, and provides the scientific community with an early opportunity to familiarize themselves with the data formats and access infrastructure for the forthcoming Legacy

Survey of Space and Time (LSST). This early release has a proprietary period of two years, during which time it is available to Rubin data rights holders only via the cloud-based Rubin Science Platform (RSP).

In this paper we have described the completion status of the observatory at the time of data acquisition, the commissioning campaign that forms the basis of DP1, and the processing pipelines used to produce early versions of data products. We provide details on the data products, their characteristics and known issues, and describe the RSP.

The data products described in this paper derive from observations obtained by LSSTComCam. LSSTComCam contains only around 5% the number of CCDs as the full LSST Science Camera (LSSTCam), yet the DP1 dataset that it has produced will already enable a very broad range of science. At 3.5 TB in size, DP1 covers a total area of ∼15 sq. deg. and contains 1792 single-epoch images, 2644 deep coadded images, 2.3 million distinct astrophysical objects, including 93 new asteroid discoveries.

While some data products expected from the LSST are not yet available, e.g. cell-based coadds, several others have been provided in DP1 that will not be available in future releases. Difference images are included in DP1, but in future releases, these will be generated on-demand via services, rather than being provided as pre-produced products. The inclusion of these images in DP1 is possible due to the small dataset size, which makes it feasible to include them at this stage. As future releases will involve much larger datasets, this approach will no longer be possible.

The RSP is continually under development, and new functionality will continue to be deployed incrementally as it becomes available, and independent of future data releases. For example, user query history capabilities, context-aware documentation and a bulk cutout services are just a few of the services currently under development.

Coincident with the release of DP1, Rubin Observatory begins its Science Validation Surveys with the LSST Science Camera. This final commissioning phase will produce a dataset that will form the foundation for the second Rubin Data Preview, DP2, expected around mid to late 2026. Full operations – marking the start of the LSST – is expected to commence by the end of 2025.

ACKNOWLEDGMENTS

This material is based upon work supported in part by the National Science Foundation through Cooperative Agreement AST-1258333 and Cooperative Support Agreement AST-1202910 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 LSSTC Institutional Members. This research has made use of NASA's Astrophysics Data System Bibliographic Services.

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

Software: Rubin Data Butler (Jenness et al. 2022), LSST Science Pipelines (Developers 2025), LSST Feature Based Scheduler v3.0 (Yoachim et al. 2024; Naghib et al. 2019) Astropy (Astropy Collaboration et al. 2013, 2018, 2022) PIFF (Jarvis et al. 2021), GBDES (Bernstein 2022), Qserv (Wang et al. 2011; Mueller et al. 2023)

APPENDIX

Glossary

Adam: Adaptive Moment Estimation. 21

ADQL: Astronomical Data Query Language (IVOA standard). 40

ADU: Analogue-to-Digital Unit. 17

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. 16

Alert: A packet of information for each source detected with signal-to-noise ratio > 5 in a difference image by Alert Production, containing measurement and characterization parameters based on the past 12 months of LSST observations plus small cutouts of the single-visit, template, and difference images, distributed via the internet. 15

Alert Production: Executing on the Prompt Processing system, the Alert Production payload processes and calibrates incoming images, performs Difference Image Analysis to identify DIASources and DIAObjects, and then packages the resulting alerts for distribution.. 20

algorithm: A computational implementation of a calculation or some method of processing. 3, 17, 19, 20, 29

AOS: Active Optics System. 3

API: Application Programming Interface. 36, 39, 40, 43

arcmin: arcminute minute of arc (unit of angle). 28

ASPIC: Analog Signal Processing Integrated Circuit. 17

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. 14, 31

ATLAS: Asteroid Terrestrial-impact Last Alert System. 15

AU: deprecated acronym for astronomical unit; use au instead. 21

au: astronomical unit. 21

B: Byte (8 bit). 29

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. 12–15, 17, 19, 20

Butler: A middleware component for persisting and retrieving image datasets (raw or processed), calibration reference data, and catalogs. 11, 13, 15, 16, 41, 42, 44

CADC: Canadian Astronomy Data Centre. 40

cadence: The sequence of pointings, visit exposures, and exposure durations performed over the course of a survey. 1

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. 4, 11, 12, 15, 16, 18, 19, 27

Camera: The LSST subsystem responsible for the 3.2-gigapixel LSST camera, which will take more than 800 panoramic images of the sky every night. SLAC leads a consortium of Department of Energy laboratories to design and build the camera sensors, optics, electronics, cryostat, filters and filter exchange mechanism, and camera control system. 1, 2

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. 9–11, 16, 19, 28

CBP: Collimated Beam Projector. 4

CCD: Charge-Coupled Device. 4, 10–12, 18, 28, 35

Center: An entity managed by AURA that is responsible for execution of a federally funded project. 14, 21, 35

Charge-Coupled Device: a particular kind of solid-state sensor for detecting optical-band photons. It is composed of a 2-D array of pixels, and one or more read-out amplifiers. 4

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).. 1, 2, 37, 38

Collimated Beam Projector: The hardware to project a field of sources onto discrete sections of the telescope optics in order to characterize spatial variations in the telescope and instrument transmission function, and to monitor filter throughput evolution during the survey. Images obtained using the CBP will be used in calibration. 4

Commissioning: A two-year phase at the end of the Construction project during which a technical team a) integrates the various technical components of the three subsystems; b) shows their compliance with ICDs and system-level requirements as detailed in the LSST Observatory System Specifications document (OSS, LSE-30); and c) performs science verification to show compliance with the survey performance specifications as detailed in the LSST Science Requirements Document (SRD, LPM-17). 1, 2

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. 3, 4, 16

CPU: Central Processing Unit. 21

CTI: Charge Transfer Inefficiency. 4, 11, 16, 17

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. 1

Data Release: The approximately annual reprocessing of all LSST data, and the installation of the resulting data products in the LSST Data Access Centers, which marks the start of the two-year proprietary period. 13, 15

Data Release Processing: Deprecated term; see Data Release Production. 17

DC2: Data Challenge 2 (DESC). 29–31

DCR: Differential Chromatic Refraction. 15, 19, 31

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. 20

deg: degree; unit of angle. 21

Department of Energy: cabinet department of the United States federal government; the DOE has assumed technical and financial responsibility for providing the LSST camera. The DOE's responsibilities are executed by a collaboration led by SLAC National Accelerator Laboratory. 1

DES: Dark Energy Survey. 15, 26

DESC: Dark Energy Science Collaboration. 29

DIA: Difference Image Analysis. 17

Difference Image Analysis: The detection and characterization of sources in the Difference Image that are above a configurable threshold, done as part of Alert Generation Pipeline. 17

Differential Chromatic Refraction: The refraction of incident light by Earth's atmosphere causes the apparent position of objects to be shifted, and the size of this shift depends on both the wavelength of the source and its airmass at the time of observation. DCR corrections are done as a part of DIA. 15, 31, 34

DIMM: Differential Image Motion Monitor. 9

Director: The person responsible for the overall conduct of the project; the LSST director is charged with ensuring that both the scientific goals and management constraints on the project are met. S/he is the principal public spokesperson for the project in all matters and represents the project to the scientific community, AURA, the member institutions of LSSTC, and the funding agencies. 45

Document: Any object (in any application supported by DocuShare or design archives such as PDMWorks or GIT) that supports project management or records milestones and deliverables of the LSST Project. 10

DOE: Department of Energy. 1

DP0: Data Preview 0. 2

DP1: Data Preview 1. 1–6, 8–19, 21, 22, 24, 26, 27, 31, 34–36, 38–41, 43–46

DP2: Data Preview 2. 2, 41, 43, 45, 46

DPDD: Data Product Definition Document. 10

DR1: Data Release 1. 45

DR3: Data Release 3. 15, 19, 27

DRP: Data Release Processing. 17

E2V: Teledyne. 4

ECDFS: Extended Chandra Deep Field-South Survey. 6, 26, 29, 31, 35, 37, 38

Education and Public Outreach: The LSST subsystem responsible for the cyberinfrastructure, user interfaces, and outreach programs necessary to connect educators, planetaria, citizen scientists, amateur astronomers, and the general public to the transformative LSST dataset. 1

EFD: Engineering and Facility Database. 44

EPO: Education and Public Outreach. 1

epoch: Sky coordinate reference frame, e.g., J2000. Alternatively refers to a single observation (usually photometric, can be multi-band) of a variable source. 2, 5, 8, 10, 13, 14, 19, 46

ESO: European Southern Observatory. 15

FBS: Feature-Based Scheduler. 5

FGCM: Forward Global Calibration Model. 19

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. 43

FITS: Flexible Image Transport System. 16

Flexible Image Transport System: an international standard in astronomy for storing images, tables, and metadata in disk files. See the IAU FITS Standard for details. 16

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. 12–15, 17, 20, 28, 29, 35

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. 14, 19, 21

FOV: field of view. 8

FrDF: French Data Facility. 17

FWHM: Full Width at Half-Maximum. 1, 3, 10, 12, 13

GAaP: Gaussian Aperture and PSF. 30, 31, 33

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. 15

Gaussian Aperture and PSF: involves Gaussianizing the PSFs and then using a Gaussian aperture (instead of top-hat) for measuring photometry. The aperture+PSF is designed to be the same across all bands, so that you measure consistent colors.. 30

HEALPix: Hierarchical Equal-Area iso-Latitude Pixelisation. 15

HiPS: Hierarchical Progressive Survey (IVOA standard). 15, 16, 41

HSM: Half-Second Moment. 24

IAU: International Astronomical Union. 15

ISR: Instrument Signal Removal. 16–18

ITL: Imaging Technology Laboratory (UA). 4, 7, 19, 25, 26, 28

IVOA: International Virtual-Observatory Alliance. 11, 13, 15, 39–41

LSST: Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope). 1–3, 5, 6, 8, 10, 13, 15, 35, 36, 44–46

LSST Science Pipelines: software used to perform the LSST data reduction pipelines.lsst.io. 4, 17, 39, 44

LSSTCam: LSST Science Camera. 2–4, 8, 27, 28, 31, 34, 35

LSSTComCam: Rubin Commissioning Camera. 2–6, 8, 10–12, 17, 19, 21, 24, 26, 35, 44–46

M1M3: Primary Mirror Tertiary Mirror. 3

M2: Secondary Mirror. 3

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). 12, 13, 16, 41

metric: A measurable quantity which may be tracked. A metric has a name, description, unit, references, and tags (which are used for grouping). A metric is a scalar by definition. See also: aggregate metric, model metric, point metric. 27

middleware: Software that acts as a bridge between other systems or software usually a database or network. Specifically in the Data Management System this refers to Butler for data access and Workflow management for distributed processing.. 17

MPC: Minor Planet Center. 35

MPCORB: Minor Planet Center Orbit database. 14, 15

National Science Foundation: primary federal agency supporting research in all fields of fundamental science and engineering; NSF selects and funds projects through competitive, merit-based review. 1

NEO: Near-Earth Object. 21

NSF: National Science Foundation. 1

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. 8, 21, 40

Operations: The 10-year period following construction and commissioning during which the LSST Observatory conducts its survey. 41

Pan-STARRS: Panoramic Survey Telescope and Rapid Response System. 15

patch: An quadrilateral sub-region of a sky tract, with a size in pixels chosen to fit easily into memory on desktop computers. 11–13, 15, 41

pipeline: A configured sequence of software tasks (Stages) to process data and generate data products. Example: Association Pipeline. 11, 17, 21, 34, 35

PNG: Portable Network Graphics. 16

POSIX: Portable Operating System Interface. 41

provenance: Information about how LSST images, Sources, and Objects were created (e.g., versions of pipelines, algorithmic components, or templates) and how to recreate them. 16

PSF: Point Spread Function. 3, 10, 12–16, 18–20, 24, 26, 35, 39

PTC: Photon Transfer Curve. 16

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. 13, 38, 40, 44

RA: Rapid Analysis. 28, 35

REB: Readout Electronics Board. 17, 18

Release: Publication of a new version of a document, software, or data product. Depending on context, releases may require approval from Project- or DM-level change control boards, and then form part of the formal project baseline. 15, 45

REST: REpresentational State Transfer. 41

RINGSS: Ring-Image Next Generation Scintillation Sensor. 9

RMS: Root-Mean-Square. 27

RSP: Rubin Science Platform. 16, 35–37, 40, 41, 43–46

Rubin Operations: operations phase of Vera C. Rubin Observatory. 19

S3: (Amazon) Simple Storage Service. 41

S3DF: SLAC Shared Scientific Data Facility. 21

schema: The definition of the metadata and linkages between datasets and metadata entities in a collection of data or archive.. 13, 44

Science Collaboration: An autonomous body of scientists interested in a particular area of science enabled by the LSST dataset, which through precursor studies, simulations, and algorithm development lays the groundwork for the large-scale science projects the LSST will enable. In addition to preparing their members to take full advantage of LSST early in its operations phase, the science collaborations have helped to define the system's science requirements, refine and promote the science case, and quality check design and development work. 29

Science Pipelines: The library of software components and the algorithms and processing pipelines assembled from them that are being developed by DM to generate science-ready data products from LSST images. The Pipelines may be executed at scale as part of LSST Prompt or Data Release processing, or pieces of them may be used in a standalone mode or executed through the Rubin Science Platform. The Science Pipelines are one component of the LSST Software Stack. 3, 26

Science Platform: A set of integrated web applications and services deployed at the LSST Data Access Centers (DACs) through which the scientific community will access, visualize, and perform next-to-the-data analysis of the LSST data products. 1, 2, 16, 37, 40, 43

SDSS: Sloan Digital Sky Survey. 15

seeing: An astronomical term for characterizing the stability of the atmosphere, as measured by the width of the point-spread function on images. The PSF width is also affected by a number of other factors, including the airmass, passband, and the telescope and camera optics. 1, 3, 10, 13, 19, 20

Sensor: A sensor is a generic term for a light-sensitive detector, such as a CCD. For LSST, sensors consist of a 2-D array of roughly 4K x 4K pixels, which are mounted on a raft in a 3x3 mosaic. Each sensor is divided into 16 channels or amplifiers. The 9 sensors that make up a raft are numbered from ”0,0” through ”2,2”. 9

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. 14, 15, 21, 26, 35

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.).. 1

sky map: A sky tessellation for LSST. The Stack includes software to define a geometric mapping from the representation of World Coordinates in input images to the LSST sky map. This tessellation is comprised of individual tracts which are, in turn, comprised of patches. 41

SLAC: SLAC National Accelerator Laboratory. 21, 25

SLAC National Accelerator Laboratory: A national laboratory funded by the US Department of Energy (DOE); SLAC leads a consortium of DOE laboratories that has assumed responsibility for providing the LSST camera. Although the Camera project manages its own schedule and budget, including contingency, the Camera team's schedule and requirements are integrated with the larger Project. The camera effort is accountable to the LSSTPO.. 21

Sloan Digital Sky Survey: is a digital survey of roughly 10,000 square degrees of sky around the north Galactic pole, plus a 300 square degree stripe along the celestial equator. 15

SLR: Stellar Locus Regression. 15

SNR: Signal to Noise Ratio. 19, 31, 35

SOAR: Southern Astrophysical Research Telescope. 9

SODA: Server-side Operations for Data Access (IVOA standard). 41

software: The programs and other operating information used by a computer.. 26, 44

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.). 20

SQL: Structured Query Language. 41

TAP: Table Access Protocol (IVOA standard). 13, 38, 40, 41

TOPCAT: Tool for OPerations on Catalogues And Tables. 43

tracklet: Links between unassociated DIASources within one night to identify moving objects. 21

tract: A portion of sky, a spherical convex polygon, within the LSST all-sky tessellation (sky map). Each tract is subdivided into sky patches. 11, 15, 16, 18, 27, 28

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. 2, 12, 14, 20

UA: University of Arizona. 4

UKDF: United Kingdom Data Facility. 17

USDF: United States Data Facility. 17, 37, 38, 44

VLT: Very Large Telescope (ESO). 15

VO: Virtual Observatory. 40

VST: VLT Survey Telescope. 15

WCS: World Coordinate System. 12, 13, 19, 26

WebDav: Web Distributed Authoring and Versioning. 41

WFD: Wide Fast Deep. 4

World Coordinate System: a mapping from image pixel coordinates to physical coordinates; in the case of images the mapping is to sky coordinates, generally in an equatorial (RA, Dec) system. The WCS is expressed in FITS file extensions as a collection of header keyword=value pairs (basically, the values of parameters for a selected functional representation of the mapping) that are specified in the FITS Standard. 12

XP: B or R Photometry (Gaia). 15

REFERENCES

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

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

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

AlSayyad, Y. 2019, Coaddition Artifact Rejection and CompareWarp, Data Management Technical Note DMTN-080, Vera C. Rubin Observatory. https://dmtn-080.lsst.io/

Ansel, J., Yang, E., He, H., et al. 2024, in 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '24) (ACM), doi: 10.1145/3620665.3640366

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., Sevilla-Noarbe, I., Drlica-Wagner, A., et al. 2025, arXiv e-prints, arXiv:2501.05739, doi: 10.48550/arXiv.2501.05739

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. M. 2022, gbdes: DECam instrumental signature fitting and processing programs, Astrophysics Source Code Library, record ascl: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

Blum, R., & the Rubin Operations Team. 2020, Vera C. Rubin Observatory Data Policy, Data Management Operations Controlled Document RDO-013, 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

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

de Vaucouleurs, G. 1948, Annales d'Astrophysique, 11, 247

—. 1953, MNRAS, 113, 134, doi: 10.1093/mnras/113.2.134

Developers, R. O. S. P. 2025, The LSST Science Pipelines Software: Optical Survey Pipeline Reduction and Analysis Environment, Project Science Technical Note PSTN-019, Vera C. Rubin Observatory, doi: 10.71929/rubin/2570545

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

Eggl, S., Juric, M., Moeyens, J., & Jones, L. 2020, in AAS/Division for Planetary Sciences Meeting Abstracts, Vol. 52, AAS/Division for Planetary Sciences Meeting Abstracts, 211.01

Esteves, J. H., Utsumi, Y., Snyder, A., et al. 2023, PASP, 135, 115003, doi: 10.1088/1538-3873/ad0a73

Euclid Collaboration, Romelli, E., Kümmel, M., et al. 2025, arXiv e-prints, arXiv:2503.15305, doi: 10.48550/arXiv.2503.15305

Fagrelius, P., & Rykoff, E. 2025, Rubin Baseline Calibration Plan, Commissioning Technical Note SITCOMTN-086, Vera C. Rubin Observatory. https://sitcomtn-086.lsst.io/

Ferguson, P., Rykoff, E., Carlin, J., Saunders, C., & Parejko, J. 2025, The Monster: A reference catalog with synthetic ugrizy-band fluxes for the Vera C. Rubin observatory, Data Management Technical Note DMTN-277, Vera C. Rubin Observatory. https://dmtn-277.lsst.io/

Fernique, P., Allen, M. G., Boch, T., et al. 2015, A&A, 578, A114, doi: 10.1051/0004-6361/201526075

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

Gray, B. 2025, find_orb: Orbit determination from observations

Guy, L. P., Bechtol, K., Bellm, E., et al. 2025, Rubin Observatory Plans for an Early Science Program, Technical Note RTN-011, Vera C. Rubin Observatory, doi: 10.5281/zenodo.5683848

Heinze, A., Eggl, S., Juric, M., et al. 2022, in AAS/Division for Planetary Sciences Meeting Abstracts, Vol. 54, AAS/Division for Planetary Sciences Meeting Abstracts, 504.04

Heinze, A., Juric, M., & Kurlander, J. 2023, heliolinx: Open Source Solar System Discovery Software

Hirata, C., & Seljak, U. 2003, MNRAS, 343, 459, doi: 10.1046/j.1365-8711.2003.06683.x

Holman, M. J., Payne, M. J., Blankley, P., Janssen, R., & Kuindersma, S. 2018, AJ, 156, 135, doi: 10.3847/1538-3881/aad69a

Homar, G. M., Tighe, R., Thomas, S., et al. 2024, in Ground-based and Airborne Telescopes X, ed. H. K. Marshall, J. Spyromilio, & T. Usuda, Vol. 13094, International Society for Optics and Photonics (SPIE), 130943C, doi: 10.1117/12.3019031

Illingworth, G., Magee, D., Bouwens, R., et al. 2016, arXiv e-prints, arXiv:1606.00841, doi: 10.48550/arXiv.1606.00841

Ingraham, P., Fagrelius, P., Stubbs, C. W., et al. 2022, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12182, Ground-based and Airborne Telescopes IX, ed. H. K. Marshall, J. Spyromilio, & T. Usuda, 121820R, doi: 10.1117/12.2630185

Ivezic, Z. 2022, Survey Cadence Optimization Committee's Phase 1 Recommendation, Project Science Technical Note PSTN-053, Vera C. Rubin Observatory. https://pstn-053.lsst.io/

Ivezić, Ž., & The LSST Science Collaboration. 2018, LSST Science Requirements Document, Project Controlled Document LPM-17, 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., et al. 2021, Mon. Not. Roy. Astron. Soc., 501, 1282, doi: 10.1093/mnras/staa3679

Jenness, T., & Dubois-Felsmann, G. P. 2025, IVOA Identifier Usage at the Rubin Observatory, Data Management Technical Note DMTN-302, Vera C. Rubin Observatory. https://dmtn-302.lsst.io/

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. 2021, Survey Strategy and Cadence Choices for the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST), Project Science Technical Note PSTN-051, Vera C. Rubin Observatory. https://pstn-051.lsst.io/

Juric, M. 2025, mpsky: Multi-purpose sky catalog cross-matching

Jurić, M., Ciardi, D., Dubois-Felsmann, G., & Guy, L. 2019, LSST Science Platform Vision Document, Systems Engineering Controlled Document LSE-319, Vera C. Rubin Observatory. https://lse-319.lsst.io/

Jurić, M., Axelrod, T., Becker, A., et al. 2023, Data Products Definition Document, Systems Engineering Controlled Document LSE-163, Vera C. Rubin Observatory. https://lse-163.lsst.io/

Kannawadi, A. 2022, Consistent galaxy colors with Gaussian-Aperture and PSF photometry, Data Management Technical Note DMTN-190, Vera C. Rubin Observatory. https://dmtn-190.lsst.io/

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

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

Lange, T., Nordby, M., Pollek, H., et al. 2024, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 13096, Ground-based and Airborne Instrumentation for Astronomy X, ed. J. J. Bryant, K. Motohara, & J. R. D. Vernet, 130961O, doi: 10.1117/12.3019302

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

Lim, K.-T. 2022, Proposal and Prototype for Prompt Processing, Data Management Technical Note DMTN-219, Vera C. Rubin Observatory. https://dmtn-219.lsst.io/

—. 2025, The Consolidated Database of Image Metadata, Data Management Technical Note DMTN-227, Vera C. Rubin Observatory. https://dmtn-227.lsst.io/

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

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

Mandelbaum, R., Hirata, C. M., Seljak, U., et al. 2005, MNRAS, 361, 1287, doi: 10.1111/j.1365-2966.2005.09282.x

Megias Homar, G., Meyers, J. M., Thomas, S. J., 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, 130943N, doi: 10.1117/12.3019361

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

Mueller, F., et al. 2023, in ASP Conf. Ser., Vol. TBD, ADASS XXXII, ed. S. Gaudet, S. Gwyn, P. Dowler, D. Bohlender, & A. Hincks (San Francisco: ASP), in press. https://dmtn-243.lsst.io

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

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

O'Mullane, W., Economou, F., Huang, F., et al. 2024a, 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

O'Mullane, W., AlSayyad, Y., Chiang, J., et al. 2024b, 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

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

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

Plazas, A. A., Shapiro, C., Smith, R., Huff, E., & Rhodes, J. 2018, Publications of the Astronomical Society of the Pacific, 130, 065004, doi: 10.1088/1538-3873/aab820

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

Plazas Malagón, A., Digel, S., Roodman, A., Broughton, A., & LSST Camera Team. 2025, LSSTCam and ComCam Focal Plane Layouts, Vera C. Rubin Observatory. https://ctn-001.lsst.io/

Refregier, A. 2003, ARA&A, 41, 645, doi: 10.1146/annurev.astro.41.111302.102207

Reiss, D. J., & Lupton, R. H. 2016, Implementation of Image Difference Decorrelation, Data Management Technical Note DMTN-021, Vera C. Rubin Observatory. https://dmtn-021.lsst.io/

Roodman, A., Rasmussen, A., Bradshaw, A., et al. 2024, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 13096, Ground-based and Airborne Instrumentation for Astronomy X, ed. J. J. Bryant, K. Motohara, & J. R. D. Vernet, 130961S, doi: 10.1117/12.3019698

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

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, Vera C. Rubin Observatory. https://dmtn-266.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

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

Swinbank, J., Axelrod, T., Becker, A., et al. 2020, Data Management Science Pipelines Design, Data Management Controlled Document LDM-151, Vera C. Rubin Observatory. https://ldm-151.lsst.io/

Taranu, D. S. 2025, The MultiProFit astronomical source modelling code, Data Management Technical Note DMTN-312, Vera C. Rubin Observatory. https://dmtn-312.lsst.io/

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

The Rubin Observatory Survey Cadence Optimization Committee. 2023, Survey Cadence Optimization Committee's Phase 2 Recommendations, Project Science Technical Note PSTN-055, Vera C. Rubin Observatory. https://pstn-055.lsst.io/

—. 2025, Survey Cadence Optimization Committee's Phase 3 Recommendations, Project Science Technical Note PSTN-056, Vera C. Rubin Observatory. https://pstn-056.lsst.io/

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

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

Waters, C. Z., Magnier, E. A., Price, P. A., et al. 2020, ApJS, 251, 4, doi: 10.3847/1538-4365/abb82b

Whitaker, K. E., Ashas, M., Illingworth, G., et al. 2019, ApJS, 244, 16, doi: 10.3847/1538-4365/ab3853

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. 2022, Survey Strategy: Rolling Cadence, Project Science Technical Note PSTN-052, Vera C. Rubin Observatory. https://pstn-052.lsst.io/

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

Zhang, T., Almoubayyed, H., Mandelbaum, R., et al. 2023, MNRAS, 520, 2328, doi: 10.1093/mnras/stac3350