

The LSST Science Pipelines Software: Optical Survey Pipeline Reduction and Analysis Environment

RUBIN OBSERVATORY SCIENCE PIPELINES DEVELOPERS¹

THE RUBIN OBSERVATORY SCIENCE PIPELINES TEAM

¹*Vera C. Rubin Observatory Project Office, 950 N. Cherry Ave., Tucson, AZ 85719, USA*

ABSTRACT

The NSF-DOE Vera C. Rubin Observatory will produce the Legacy Survey of Space and Time (LSST), producing 11 data releases over the ten-year survey. The LSST Science Pipelines Software, the Optical Survey Pipeline Reduction and Analysis Environment (OSPRAE), will be used to create these data releases and to perform the nightly alert production. This paper provides an overview of the LSST Science Pipelines Software, describing the components and how they are combined to form pipelines.

Keywords: Astrophysics - Instrumentation and Methods for Astrophysics — methods: data analysis — methods: miscellaneous

1. INTRODUCTION

The NSF-DOE Vera C. Rubin Observatory will be performing the 10-year Legacy Survey of Space and Time (LSST; [Ž. Ivezić et al. 2019](#)) starting in 2025. Rubin Observatory is located on Cerro Pachon in Chile and consists of the 8.4m Simonyi Survey Telescope ([S. J. Thomas et al. 2022](#)) with the 3.2-gigapixel LSSTCam survey camera ([SLAC National Accelerator Laboratory & NSF-DOE Vera C. Rubin Observatory 2025](#); [A. Roodman et al. 2024](#)) performing the main survey and the Rubin Auxiliary Telescope ([P. Ingraham et al. 2020](#)) providing supplementary atmospheric calibration data using the LATISS instrument ([NSF-DOE Vera C. Rubin Observatory 2020](#)). The Data Management System (DMS; [W. O’Mullane et al. 2022](#)) is designed to handle the flow of data from the telescope, approaching 20 TB per night, in order to issue alerts and to prepare annual data releases. A central component of the DMS is the LSST Science Pipelines software that provides the algorithms and frameworks required to process the data from the LSST and generate the coadds, difference images, and catalogs to the user community for scientific analysis.

The LSST Science Pipelines software consists of the building blocks and pipeline infrastructure required to construct high performance pipelines to process the data from the LSST. It has been under development since at least 2004 ([T. Axelrod et al. 2004](#)) and has evolved sig-

nificantly over the years as the project transitioned from prototyping ([T. Axelrod et al. 2010](#)) and entered into formal construction ([M. Jurić et al. 2017](#)). The software is designed to be usable by other optical telescopes and this has been demonstrated with Hyper Suprime-Cam on the Subaru Telescope in Hawaii ([J. Bosch et al. 2018](#)) and also with data from the Dark Energy Camera (DECam), the VISTA infrared camera (VIRCAM), the Wide Field Survey Telescope (WFST; [M. Cai et al. 2025](#)), and the Gravitational-wave Optical Transient Observer (GOTO; [J. R. Mullaney et al. 2021](#)). It was used to make the first public data release from Rubin Observatory, the LSST Data Preview 1 ([Vera C. Rubin Observatory 2025](#); [NSF-DOE Vera C. Rubin Observatory 2025](#)), which consisted of data from the LSST Commissioning Camera (LSSTComCam; [SLAC National Accelerator Laboratory & NSF-DOE Vera C. Rubin Observatory 2024](#); [B. Stalder et al. 2022](#)).

In this paper we provide an overview of the software system, dividing it into four rough tiers:

- low-level utility code and our data access, configuration, and execution frameworks, and our core algorithmic primitives are described in sections 2, 3, and 4, respectively;
- reusable mid-level algorithmic components are described in 5;

- high-level tasks and pipelines are described in 6, along with some of the details of the algorithms specific to them;
- analysis and validation tooling is described in 7.

There are no sharp boundaries between these tiers; they are better considered to be a loose organizational aid than any kind of formal classification. We do not include details of the science validation of the individual algorithms, and do not attempt to cover components of the system in uniform detail; instead we focus attention on those that are unique or novel or not well described elsewhere. The other components of the LSST DMS, such as the workflow system (M. Gower et al. 2022; E. Karavakis et al. 2024), the Qserv database (D. L. Wang et al. 2011; F. Mueller et al. 2023) and the Rubin Science Platform (M. Jurić et al. 2019; W. O’Mullane et al. 2024), are not covered in this paper.

2. FUNDAMENTALS

The LSST Science Pipelines software is written in Python with C++ used for high-performance algorithms and for core classes that are usable in both languages. We use Python 3 (having ported from python 2, T. Jenness 2020, currently with a minimum version of Python 3.12), and the C++ layer can use C++17 features with pybind11 being used to provide the interface from Python to C++. Additionally, the C++ layer uses `ndarray` to allow seamless passing of C++ arrays to and from Python `numpy` arrays. This compatibility with `numpy` is important in that it makes LSST data structures available to standard Python libraries such as Scipy and Astropy (T. Jenness et al. 2016; Astropy Collaboration et al. 2018).

Although all the software uses the `lsst` namespace, the code base is split into individual Python products in the LSST GitHub organization² that can be installed independently and which declare their own dependencies. These dependencies are managed using the “Extended Unix Product System” (EUPS; N. Padmanabhan et al. 2015; T. Jenness et al. 2018) where most of the products are built using the SCons system (S. Knight 2005) with LSST-specific extensions provided in the `sconsUtils` package enforcing standard build rules and creating the necessary Python package metadata files.

For logging we always use standard Python logging with an additional `VERBOSE` log level between `INFO` and `DEBUG` to provide additional non-debugging detail that can be enabled during batch processing. This verbose




Figure 1. The number of lines of code comprising the LSST Science Pipelines software as a function of year. Line counts include comments but not blank lines. Python interfaces are implemented using `pybind11` and that is counted as C++ code. For the purposes of this count Science pipelines software is defined as the `lsst_distrib` metapackage and does not include code from third party packages or code from the solar system pipeline.

logging is used for periodic logging where long-lived analysis tasks are required to issue a log message every 10 minutes to indicate to the batch system that they are still alive and actively performing work. For logging from C++ we use `Log4CXX` wrapped in the `lsst.log` package to make it look more like standard Python logging, whilst also supporting deferred string formatting such that log messages are only formed if the log message level is sufficient for the message to be logged. These C++ log messages are forwarded to Python rather than being issued from an independent logging stream. Finally, we also provide some LSST-specific exceptions that can be thrown from C++ code and caught in Python.

As of April 2025, the Science Pipelines software is approximately 700,000 lines of Python and 225,000 lines of C++. The number of lines in the pipelines code as a function of time is given in Fig. 1.

2.1. Python environment

An important aspect of running a large data processing campaign is to ensure that the software environment is well defined. We define a base python environment using `conda-forge` via a meta package named `rubin-env`³. This specifies all the software needed to build and run the science pipelines software. A Docker container is built for each software release and the fully-specified versions of all software are recorded to ensure repeatability.

² <https://github.com/lsst>

³ <https://github.com/conda-forge/rubinenvironment-feedstock>

2.2. Unit Testing and Code Coverage

Unit testing and code coverage are critical components of code quality (T. Jenness et al. 2018). Every package comes with unit tests written using the standard `unittest` module. We run the tests using `pytest` (H. Krekel 2017) and this comes with many advantages in that all the tests run in the same process and requiring global parameters to be well understood, tests can be run in parallel in multiple processes, plugins can be enabled to extend testing and record test coverage, and a test report can be created giving details of run times and test failures. Coding standards compliance with PEP 8 (G. van Rossum 2013) is enforced using GitHub Actions, the `ruff` package, and `pre-commit` checks.

A Jenkins system provides the team with continuous integration across multiple packages. This includes longer tests (up to a few hours) in which we run complete pipelines on small precursor datasets (typically a few GB) fetched via `git-lfs`.

3. DATA ACCESS AND EXECUTION ABSTRACTIONS

The algorithmic components of the LSST Science Pipelines are built on a suite of packages that together form a powerful data access and execution framework (`pex_config`, `resources`, `daf_butler`, `pipe_base`, `ctrl_mpexec`, and `ctrl_bps`). Unlike most of the rest of the codebase, these packages can be individually installed with `pip` as well as EUPS and can be used on their own.

3.1. Butler

Early in the development of the LSST Science Pipelines software it was decided that the algorithmic code should be written without knowing where files came from, what format they were written in, where the outputs are going to be written or how they are going to be stored. All that the algorithmic code needs to know is the relevant data model and the Python type. To meet these requirements we developed a library called the Data Butler (see e.g., T. Jenness et al. 2022; N. B. Lust et al. 2023).

The Butler internally is implemented as a registry, a database keeping track of datasets, and a datastore, a storage system that can map a Butler dataset to a specific collection of bytes. A datastore is usually a file store (including POSIX file system, S3 object stores, or WebDAV) but it is also possible to store metrics directly into the Sasquatch metrics service (A. Fausti 2023; A. Fausti Neto et al. 2024).

A core concept of the Butler is that every dataset must be given what we call a “data coordinate.” The

Table 1. Common dimensions present in the default dimension universe.

Name	Description
<code>instrument</code>	Instrument.
<code>band</code>	Waveband of interest.
<code>physical_filter</code>	Filter used for the exposure.
<code>day_obs</code>	The observing day.
<code>group</code>	Group identifier.
<code>exposure</code>	Individual exposure.
<code>visit</code>	Collection of 1 or 2 exposures.
<code>tract</code>	Tessellation of the sky.
<code>patch</code>	Patch within a tract.

data coordinate locates the dataset in the dimensional space where dimensions are defined in terms that scientists understand. Some commonly used dimensions are listed in Table 1. Each dataset is uniquely located by specifying its dataset type, its run collection, and its coordinates, with Butler refusing to accept another dataset that matches all three of those values. The dataset type defines the relevant dimensions (such as whether this is referring to observations or a sky map) and the associated Python type representing the dataset. The run collection can be thought of as a folder grouping datasets created by the same batch operation, but does not have to be a folder within a file system.

As a concrete example, the file from one detector of an LSSTCam observation taken sometime in 2025 could have a data coordinate of `instrument="LSSTCam", detector=42, exposure=2025080300100` and be associated with a `raw` dataset type. The `exposure` record itself implies other information such as the physical filter and the time of observation. A deep coadd on a patch of sky would not have `exposure` dimensions at all and would instead be something like `instrument="LSSTCam", tract=105, patch=2, band="r", skymap="something"`, which would tell you exactly where it is located in the sky and in what waveband since you can calculate it from the tract, patch, band and skymap.

3.2. Pipelines and Tasks

The data dimensions system also plays a fundamental role in how the LSST processing pipelines are assembled and run; high-level pieces of algorithmic code called `PipelineTasks` declare the dimensions of their units of work (“quanta”), their inputs, and their outputs, allow-

ing a directed acyclic graph (a “quantum graph”) describing the processing to be assembled from a YAML declaration of the tasks to be run, their configuration, and a butler database query. Quantum graphs can range in size from a few tens of quanta (e.g., for the nightly processing performed on a single detector image) to millions (for a piece of the yearly data release pipelines), and serve as the common interface for multiple execution systems, including the low-latency nightly Prompt Processing framework and the Batch Processing System (BPS; M. Gower et al. 2022), which adapts quantum graphs for execution at scale by third-party workflow management systems like HTCondor (HTCondor Team 2024), Parsl (Y. Babuji et al. 2019), and PanDA (E. Karavakis et al. 2024).

TODO: examples of pipeline YAML, pipeline graph diagrams

Algorithmic code below the `PipelineTask` level is often subdivided into multiple “subtasks” that (like `PipelineTask` itself) inherit from the base `Task` class, which provides easy access to hierarchical logging, meta-data, and configuration.

3.3. Pipeline Visualization

Visualizing pipeline execution is crucial for understanding task dependencies, debugging, optimizing workflows, and ensuring correct data flow within the LSST Science Pipelines. To support this, `pipetask` build provides several options for visualizing the pipeline graph—a simplified directed acyclic graph that shows how tasks relate to dataset types, without including data IDs.

A text-based view can be generated using `pipetask` build `--show pipeline-graph`, which outputs an ASCII-style diagram. This format is especially useful for quick inspection or when working in a terminal-only environment. For graphical visualization, the `--pipeline-dot` and `--pipeline-mermaid` options export the pipeline graph in Graphviz DOT⁴ and Mermaid⁵ formats, respectively. The Mermaid format is particularly well-suited for sharing in accessible, web-based contexts.

Unlike DOT files, which typically require rendering with external tools like Graphviz’s `dot`, Mermaid definitions can be directly rendered in Markdown-based platforms such as GitHub, GitLab, some Jupyter environments, and even Slack with the appropriate plugin. This makes Mermaid an effective format for generating interactive, easily shareable pipeline graphs that can be di-

rectly embedded in documentation, notebooks, or code review tools.

Figure 2 shows a visualization of a subset of two tasks from the `LSSTComCam/DRP-v2.yaml` pipeline using the Mermaid format. The diagram shows the relationships between tasks and their input and output datasets as well as the sequence in which the tasks are expected to run. Such visualizations can help uncover misconfigurations, missing inputs, or unexpected data dependencies that might otherwise result in issues such as empty QuantumGraphs or failed pipeline execution.

3.4. Configuration

Pex Config is the foundational configuration system for the LSST Rubin Observatory’s ambitious science pipelines. It’s far more than a simple parameter parser; it’s a framework that mediates between diverse configuration sources and the complex software that processes astronomical data. At its core, Pex Config functions as an intermediate representation, decoupling the pipelines from the specifics of configuration file formats (like YAML, JSON) and providing a unified, Python-native interface to all configurable parameters. This intermediate representation, resembling a Domain Specific Language embedded within Python, also allows leveraging the full power of a programming language for parsing or setting configuration values. An example of this can be seen in the following code block which shows a fragment used to configure one of the shape measurement routines. This abstraction is critical for maintainability, allowing the underlying file formats and or execution systems to evolve without impacting the pipeline code. It also provides a mechanism to deprecate configurables which will change in future versions of the software stack, allowing users an easy migration path.

```

1 import os.path
2 from lsst.utils import getPackageDir
3
4 try:
5     location =
6         getPackageDir("meas_extensions_shapeHSM")
7     path = os.path.join(, "config", "enable.py")
8     config.load(path)
9     plugins = config.plugins
10    plugin =
11        plugins["ext_shapeHSM_HsmShapeRegauss"]
12    plugin.deblendNChild = "deblend_nChild"
13    # Enable debiased moments
14    config.plugins.names |=
15        ["ext_shapeHSM_HsmPsfMomentsDebiased"]
16 except LookupError as e:
17     print("Cannot enable shapeHSM (%s): disabling
18         HSM shape measurements" % (e,))

```

Listing 1. Code configuration in python

⁴ <https://www.graphviz.org>

⁵ <https://mermaid.js.org>




Figure 2. Example pipeline visualization of four selected tasks from the `LSSTComCam/DRP-v2.yaml` pipeline in the Mermaid format. The diagram illustrates the flow of datasets between tasks, with dashed lines indicating prerequisite inputs. This visualization helps validate task dependencies and the expected sequence of execution.

The design of Pex Config centers around the concepts of “Fields” and “Config” objects. Fields represent individual configurable values – things like exposure times, image quality thresholds, or database connection strings. Each Field is strongly typed, supporting a variety of data types (such as integers, floats, strings, booleans, and lists). Config objects, on the other hand, are containers that group related Fields together, creating logical units of configuration. One of the highlights of Pex Config is its composability. Config objects can be nested within other Config objects using a special “ConfigField,” allowing for the creation of complex, hierarchical configuration trees that mirror the structure of the pipelines themselves. This allows for modularity and reuse of configuration components across different parts of the system.

A strength of Pex Config is its flexible application of configuration values. Values can be set at multiple stages: via command-line arguments, loaded from configuration files, or defined directly within the pipeline code. Importantly, these stages are applied progressively, with later stages overriding earlier ones. This allows for a powerful combination of default settings, user-defined customizations, and dynamic adjustments. Mechanisms also exist to apply values to all instances of a particular Config object within a tree, simplifying the management of shared parameters and ensuring consistency.

Beyond runtime configuration, Pex Config is deeply concerned with data provenance and reproducibility. It provides mechanisms for persisting and restoring configuration values, allowing for complete tracking of pipeline parameters used in a particular data processing run. Crucially, it also maintains a history of each Field’s value, recording when and where it was set – whether via

the command line, a configuration file, or programmatically. This detailed history is invaluable for debugging, auditing, and ensuring the reproducibility of scientific results. The system also incorporates robust validation mechanisms, enabling checks on individual Fields and groups of values before they are used by the pipelines, preventing errors and ensuring data quality. Validation can range from simple type checking, ensuring values fall within acceptable ranges or specific patterns, to complex custom functions that enforce specific constraints.

Finally, Pex Config is designed with documentation in mind. All Fields and Config objects can be richly documented using documentation strings and attributes. This documentation structure is not only readable by humans but can also be parsed by automated tools to generate comprehensive documentation pages, eliminating the need for manual documentation creation. This ensures that the configuration system is well-documented and easy to understand, even for new developers. The system is flexible enough that it has been adopted by the DRAGONS software (K. Labrie et al. 2023).

3.5. Instrument Abstractions: Obs Packages

The Butler and pipeline construction code know nothing about the specifics of a particular instrument. In the default dimension universe there is an `instrument` dimension that includes a field containing the full name of a Python `Instrument` class. This class, which uses a standard interface, is used by the system to isolate the instrument-specific from the pipeline-generic. Some of the responsibilities are:

- Register instrument-specific dimensions such as `detector`, `physical_filter` and the default `visit_system`.

- Define the default `raw` dataset type and the associated dimensions.
- Provide configuration defaults for pipeline task code that is processing data from this instrument.
- Provide a “formatter” class that knows how to read raw data.
- Define the default curated calibrations known to this instrument.

The `Instrument` interface is defined in two levels: the minimal interface in the `pipe_base` package defines everything needed to use the butler and execution system, while a more complete subclass in `obs_base` provides considerable additional functionality but is not in the minimal, `pip`-installable suite.

By convention we define the instrument class and associated configuration in `obs` packages. There are currently project-supported `obs` packages for:

- LSSTCam (SLAC National Accelerator Laboratory & NSF-DOE Vera C. Rubin Observatory 2025; A. Roodman et al. 2024; S. M. Kahn et al. 2010), LATISS (NSF-DOE Vera C. Rubin Observatory 2020; P. Ingraham et al. 2020), and associated Rubin Observatory test stands and simulators.
- Hyper-SuprimeCam on the Subaru telescope (S. Miyazaki et al. 2018).
- The Dark Energy Camera on the CTIO Blanco telescope (B. Flaugher et al. 2015; D. L. DePoy et al. 2008).
- CFHT’s MegaPrime (O. Boulade et al. 2003).

Additionally, teams outside the project have developed `obs` packages to support Subaru’s Prime Focus Spectrograph (S.-Y. Wang et al. 2020), VISTA’s VIR-CAM (W. Sutherland et al. 2015), the Wide Field Survey Telescope (WFST; M. Cai et al. 2025), and the Gravitational-wave Optical Transient Observer (GOTO; J. R. Mullaney et al. 2021).

3.6. Metadata Translation

Every instrument uses different metadata standards but the Butler data model and pipelines require some form of standardization to determine values such as the coordinates of an observation, the observation type, or the time of observation. To perform that standard extraction of metadata each supported instrument must provide a metadata translator class using

the `astro_metadata_translator` infrastructure.⁶ The translator classes can understand evolving data models and allow the standardized metadata to be extracted for the lifetime of an instrument even if headers changed. Furthermore, in addition to providing standardized metadata the package can also provide programmatic or per-exposure corrections to data headers prior to calculating the translated metadata. This allows files that were written with incorrect headers to be recovered during file ingestion.

4. CORE ALGORITHMIC PRIMITIVES AND DATA STRUCTURES

The high-level algorithms in the LSST Science Pipelines are largely built from algorithmic primitives and data structures implemented in the `afw` package. `afw` is a large, complex, C++-heavy suite of multiple libraries that sometimes suffer from historical idiosyncrasies, but are nevertheless extremely powerful and well optimized. These include (but are not limited to):

- `afw.geom` holds our high-level geometry primitives. This includes composable coordinate transforms and world coordinate systems (WCS). It also includes `SpanSet`, a run-length encoding (RLE) description of a set of pixels in a 2-d image, a simple `Polygon` class, and routines for working with various ellipse parameterizations.
- `afw.math` includes convolution, resampling, general-purpose interpolation, statistics, and least-squares fitting algorithms.
- `afw.detection` contains threshold-based detection on images and point-spread function (PSF) model interfaces. This includes the `Footprint` class, which combines a `SpanSet` with a list of peaks to represent either a single source detection or a group of blended sources.
- `afw.image` centers around the `Exposure` class, which combines an image (`Image`), bitmask (`Mask`), and variance image with the many objects used to astrophysically characterize an observation or coadd (PSF, WCS, aperture corrections, etc).
- `afw.table` holds data structures for tabular data, with both row- and column-based views. The `afw.table.io` package defines a framework for persisting arbitrary objects to a series of FITS binary table HDUs, which is used in the on-disk form of the `Exposure` class.

⁶ <https://astro-metadata-translator.lsst.io>

- `afw.cameraGeom` provides a hierarchical description of large-format photometric cameras (like LSSTCam, HSC, or DECam), including optical distortions, focal plane layouts, and amplifier regions.

Most of the algorithms and data structures in `afw` are implemented there directly, and often represent evolutions of concepts first developed in the SDSS *Photo* pipeline or the Pan-STARRS Image Processing Pipelines. Others delegate to third-party libraries, particularly Eigen (linear algebra), the Gnu Scientific Library (interpolation, random numbers), Boost (image iterators and geometry), CFITSIO (image and table I/O) and Starlink AST (coordinate systems and transforms; [D. S. Berry et al. 2016](#)).

Even lower-level data structures are defined in a handful of packages just below `afw`. The `sphgeom` package is used for spherical geometry calculations, sky-based regions, and hierarchical sky pixelization schemes, while `geom` provides simple 2-d Euclidean `Point`, `Extent`, and `Box` types in both integer- and floating-point variants. `geom` also includes linear transforms and (for historical reasons) its own angle-manipulation and sky coordinate type.

C++ mapping types (with Python bindings) and date/time objects are defined in `daf_base`. The `DateTime` package is used in our C++ data structures mostly to represent TAI times. The `PropertySet` represents a hierarchical key/value data structure whereas `PropertyList` is a flat data structure that is used to represent a FITS header and supports multi-valued keys and key comments.

Another small set of core packages sits just above `afw`:

- `skymap` defines interfaces and a few implementations of our system for mapping the sky onto a set of slightly-overlapping image-friendly projections and tiles for coaddition. Each distinct projection in a skymap is called a `tract`, and each `tract` is further divided into multiple `patches`.
- `shapelet` includes optimized evaluation of Gauss-Hermite and Gauss-Laguerre functions and their derivatives.

A substantial fraction of our core packages predate the now-ubiquitous Astropy package, and in some cases we now prefer to use Astropy types in *most* new code (date/time representations and tables in particular) and especially public interfaces, following the recommendations from [T. Jenness et al. \(2016\)](#). Fully retiring core libraries that have Astropy counterparts is at best a long-term project, however, due to our continued need

for these objects in considerable amounts of C++ that has no equivalent in Astropy (or anywhere else).

5. KEY ALGORITHMIC COMPONENTS

Most of the key algorithms we have implemented for processing Rubin data are used in multiple pipelines, and in many cases an algorithmic component is used multiple times within a single pipeline. For the most part, these reusable algorithms are implemented as regular `Task` objects, and these are combined in higher-level `PipelineTasks` discussed later in [6](#). In other cases, the algorithm is inseparable from its I/O, and is implemented directly as a `PipelineTask`. This section includes examples of both.

5.1. Instrument Signature Removal

Raw images from charge-coupled devices (CCDs) contain instrumental effects, such as dark currents, clocking artifacts, or crosstalk between neighboring amplifiers, that can be removed in the data processing. In the Rubin pipeline, this step is called Instrument Signature Removal (ISR) and is the first processing applied to a raw CCD exposure. The package performing the ISR on an exposure, called `ip_isr`, is detailed below in [§5.1.1](#): it is a critical package for Data Release Production (DRP) pipelines used to process LSST images and requires calibration products produced and verified by `cp_pipe` and `cp_verify` respectively as described in [§6.5](#). For further information about the life cycle of a calibration product and the procedures it entails, see [C. Waters & E. Rykoff \(2025\)](#). A general overview of the ISR steps (based on the model in [Fig. 3](#)) and calibration products production (including generation, verification, certification, approval, and distribution) is given in [A. A. Plazas Malagón et al. \(2025\)](#).

We note that we focus here on our approach to performing ISR on data from LSST cameras only (LSSTCam, LSSTComCam, and LATISS), although we also provide calibration pipelines for other cameras such as DECam and HSC using a different ISR approach.

5.1.1. ISR package

Exposures from LSST cameras are affected by instrumental effects, ranging from well-known CCD effects like dark currents or bias levels to effects more recently characterized like tree-rings (see [H. Y. Park et al. \(2017\)](#); [H. Park et al. \(2020\)](#); [J. H. Esteves et al. \(2023\)](#); [Y. Okura et al. \(2015, 2016\)](#) for more details on tree rings in LSSTCam and their impact on science) or the Brighter-Fatter effect as discussed in [A. Broughton et al. \(2024\)](#). Correcting for these effects requires specific calibrations, which we refer to as calibration products. In LSST cameras, calibration products typically are a combined bias,




Figure 3. Schematic of the instrument model for detector effects in LSST cameras which `isrTaskLSST` is based on at the time of publication. More details about the model can be found in P. Fagrelus & E. Rykoff (2025) and A. A. Plazas Malagón et al. (2025).

a combined dark, a Photon Transfer Curve (PTC), a crosstalk matrix, a list of defects, and a look-up table of non-linearity parameters. The meaning of these calibration products and the details on the Rubin Observatory’s ISR and calibration approach can be found in A. A. Plazas Malagón et al. (2025) and P. Fagrelus & E. Rykoff (2025).

The `ip_isr` package⁷ contains the codes needed to remove instrument signatures in exposures from LSST cameras and to produce calibration products. To inform our ISR approach, we first designed a model of the instrument, displayed in Fig. 3, based on our knowledge of the hardware and electronics. This model states the order in which the different known instrumental effects happen, from a photon hitting the CCD to the output ADC unit (ADU) signal. In turn, `isrTaskLSST` in `ip_isr` sequentially applies corrections of these effects in the opposite order as their effects occur in the model, as we are attempting to remove the impact of those effects on the image. Such corrections are typically done by calling other `Tasks` (e.g. overscan, crosstalk, etc.) also implemented in `ip_isr`. Not shown in Fig. 3 is a final amplifier offset (amp-offset) correction step. The amp-offset correction is applied to the output exposure at the end of `isrTaskLSST` to remove discontinuities in the background sky level across amplifier boundaries. For more details on the amp-offset correction, see Appendix §A.

Overall, `isrTaskLSST` takes a raw CCD exposure, and calibration products if available, and outputs a `Struct` containing the output exposure, the `post_isr_image`

output exposure as well as its binned version for easier display, the exposure without interpolation and statistics on the output exposure. `IsrTaskLSSTConfig` defines the configurations used in this `Task`, with them set by default to their expected values to perform ISR on a typical LSSTCam exposure. Configuration parameters starting with `do` will typically correspond to an ISR step. These steps are turned on or off in the pipelines when producing the different calibration products. We have also developed `isrMockLSST` which simulates a raw exposure and corresponding calibration products and is used to test `isrTaskLSST`.

5.2. Background Subtraction

Our low-level background subtraction algorithms (in the `meas_algorithms` package) operate by binning images, masking out detections and bad pixels, and computing robust statistics (clipped means, by default) on the remaining values in those bins. The bin averages can then be interpolated via Akima splines or approximated by Chebyshev polynomials.

This code is run and re-run in many different configurations in many different contexts.

- It is the first step run after Instrument Signature Removal (§5.1.1) on each detector in both the nightly and data release pipelines.
- In our most advanced configurations of the data release pipelines, we replace these original per-detector backgrounds with new backgrounds fit across the full visit (H. Aihara et al. 2019). These are configured to subtract the background on large scales, to preserve low surface-brightness features as much as possible.
- We sometimes run another round of background subtraction on coadds before detecting objects on them, especially when only full-visit backgrounds were subtracted from the input images. Unfortunately we have found that running background subtraction with a fairly small spatial scale (e.g. 256-pixel bins) – which intentionally oversubtracts bright stars and low surface-brightness features – is necessary for the best photometry on most objects, but we prefer to subtract the small-scale background as late as possible in order to allow it to be easily restored for other science cases.

In addition, we routinely run background subtraction as a small tweak every time we run detection (§5.3). We *define* our background to be the sum of all light that we do not consider part of a detection, which means it needs to be adjusted when the detection threshold changes (e.g. when we detect fainter objects on coadds).

⁷ https://github.com/lsst/ip_isr

5.3. Source Detection

In the LSST Science Pipelines, *detection* refers specifically to the process of finding above-threshold regions in images and one or more peaks within them (i.e. **Footprints**, as described in 4). Before this thresholding, we convolve each image with an approximation of its PSF model, as this is (approximately) optimal for detecting isolated point sources (J. Bosch et al. 2018). After detection, **Footprints** with multiple peaks are generally deblended (5.4) before being measured (5.5). Our detection algorithms assume the image has already been background-subtracted.

All of our detection tasks are implemented in the `meas_algorithms` package.

5.3.1. SourceDetectionTask

SourceDetectionTask convolves the image with a Gaussian approximation to the exposure PSF and detects **Footprints** above a configurable threshold in either signal-to-noise or absolute flux level. It can optionally detect negative deviations as well as positive (which is useful when operating on difference images).

A serious challenge for this algorithm in deep or otherwise crowded fields is that noise peaks or below-threshold sources can be pushed above the detection threshold when they land in the wings of brighter neighbors. To reduce the number of spurious peaks due to this effect, **SourceDetectionTask** can optionally perform a *temporary* small-scale background subtraction before looking for peaks (but after finding the above-threshold regions); while the spline and Chebyshev models are not good models for the wings of large objects, they are often better than nothing.

5.3.2. DynamicDetectionTask

The **DynamicDetectionTask** is a specialized version of **SourceDetectionTask** that adjusts detection thresholds based on the local background and noise. This task was initially developed to address detection efficiency issues noted in HSC data, though to be related to correlations in the noise due to warping. First, **DynamicDetectionTask** detects sources using a lower detection threshold than normal. In so doing, we identify regions of the sky which are unlikely to contain real source flux. Next, a configurable number of sky objects (see `refsec:sky-objects`) are placed in these sky regions (1000 by default), and the PSF flux and standard deviation for each of these measurements is calculated. Using this information, we set the detection threshold such that the standard deviation of the measurements matches the median estimated error.

5.3.3. MaskStreaksTask

Instead looking for astrophysical sources like our other detection tasks, **MaskStreaksTask** is responsible for identifying and masking pixels affected by streaks (primarily artificial satellites). It searches for linear features with a Canny filter and the Kernel-Based Hough Transform (L. A. F. Fernandes & M. M. Oliveira 2008), and operates most effectively on difference images. The algorithm is described in more detail in C. Saunders (2021), which focuses on its usage on differences between PSF-matched warps and a PSF-matched coadd, as described in 6.2.1. This task is also used in traditional difference images, i.e. those formed by subtracting a template coadd from a science images, as described in 6.3.

5.4. Deblending

Deblending has become the standard astronomical term for dealing with images where multiple distinct astrophysical sources overlap. In the LSST Science Pipelines, it specifically means assigning different fractions of the flux of all of the pixels in a **Footprint** to each of the peaks in that **Footprint**, which are then each considered a “child” source. This fractional flux assignment is very different from creating a “segmentation map” that fully assigns each pixel to each child source, as in e.g. Source Extractor (E. Bertin & S. Arnouts 1996). **HeavyFootprint** (an extension of **Footprint** that adds a flattened array of pixel values) is used to pass the per-pixel fluxes to downstream algorithms.

Deblending in the science pipelines is performed differently for single-band (visit) image processing vs. multi-band (coadd) image processing. This section gives a basic description of each algorithm.

5.4.1. Single-band Deblending

Deblending on single-band images (ie. visit) is performed using the `meas_deblender` package and is based on the deblender used in SDSS (citation needed), with a few differences that will be discussed shortly. Similar to the SDSS deblender, the LSST deblender creates a template for each source in a blend using a very simple (yet computationally efficient) model for each peak position in a parent **Footprint**. Once a template has been created for each peak in the blend, the deblender combines all of the source templates into a single blend model by summing their values in each pixel. For each pixel in a source template, the ratio of the source template value to the total blend model is calculated and used to weight the pixel value from the image to create a model for each source. The source models are thus flux conserving in that adding them together will yield the original image except for pixels that do not appear

in any of the individual templates. A cleanup algorithm is then run to allocate the remaining pixels to one of the sources in the blend based on a set of criteria including distance to the center, brightness of the nearest sources, etc.

5.4.2. Deblender Template Generation

The main ansatz of the SDSS deblending algorithm is that the flux from stars and galaxies in a ground based telescope is nearly 180 degree symmetric. Figure 4 illustrates how a 1D slice through the center of two blended sources can exploit this symmetry by setting the pixels on opposite sides of the (integer) center pixel to the minimum value of both pixels. In other words, for simple blends of only two sources the deblender can use the flux on the non-blended side to constrain the value of the flux on the blended side. Despite the fact that stars (PSFs) and galaxies are not exactly symmetric, especially since their position is not exactly centered in the center of a single pixel, this algorithm works quite well for generating templates in simple blends that very nearly approximate each source when redistributing flux from the image.

For sources with low SNR the algorithm fails due to noise in the image, generating galaxy templates that are typically very jagged and unphysical. To combat this, for each **peak** in the parent **Footprint** the deblender first attempts to fit the flux from the image with a simple PSF model that allows its position, amplitude, and a linear background, to vary. If the fit has a reasonable χ^2 value then the deblender will use this scaled PSF model as a template for the source. Only for sources that cannot be adequately modeled with the PSF are the symmetric templates used.

The main failure point of this algorithm is when three (or more) sources lie along the same axis. For example, Figure 5 illustrates a 1D slice through the center of three aligned sources. In this case the minimum pixel on each side of the central source cannot constrain the flux at that radial location and results in a template that has extra bumps from its neighbors. This turns out to be more catastrophic than one might expect. Notice that even the neighboring sources, which have very good templates created by using symmetry on their unblended side, have their resulting models contaminated due to the central object that steals flux from both of them. In single visits the number of "three in a row" blends is small enough that we sacrifice the quality of the models for efficiency and still use the single-band deblender. For LSST-depth coadds this becomes a significant problem, as deep coadds can have as much as




Figure 4. A 1D slice of two blended Gaussian sources illustrating how symmetry can be utilized to model blended sources.




Figure 5. A 1D slice through three aligned Gaussian sources, demonstrating a failure case of relying on symmetry for generating deblender templates. Notice that for sources 2 and 3 the templates are reasonable but due to the inability of source 1 to use symmetry to constrain flux in the blended region, the resulting models for all three sources are poor. This catastrophic "three in a row" problem was part of the motivation for creating **scarlet** to incorporate spectral information and a more rigorous iterative deblending algorithm.

40% of blends having 3 or more sources and a more sophisticated algorithm is needed.

5.4.3. Multi-band Deblending

The multi-band deblender is an implementation of the scarlet deblending algorithm described in [P. Melchior et al. \(2018\)](#). In our implementation, **scarlet_lite**, we have made our own set of simplifying assumptions that are different from the original scarlet algorithm to make it more efficient when used in a large ground based survey like LSST. Similar to the original **scarlet** we

make the assumption that astrophysical objects can be thought of as a collection of components, where each component has the properties

- Components have a single color (spectrum) that is the same in all pixels over its shape (morphology)
- Components have flux that monotonically decreases from the center
- Component flux is additive

The classic example is decomposing a single galaxy into bulge and disk components, where both the bulge and disk share a common center but have different spectra and morphologies. Something more complex, like a grand design spiral, could in theory be modeled as a source with multiple components, where spiral arms and star forming regions could still be thought of as separate monotonic components. For the science pipelines we ignore those more complicated structures, as detection typically already shreds large galaxies into multiple sources. Instead we use a signal to noise cut where low flux sources are modeled with a single component and higher flux sources are modeled with two components.

Scarlet lite initializes models with nearly the same templates as those generated by the single-band deblender. Using a χ^2 -like monochromatic image created by weighting each band by its inverse variance, scarlet lite creates initial morphology models that are symmetric from the center in the monochromatic image, with the additional constraint that the flux is monotonically decreasing from the center. In order to satisfy the constraint that all pixels in the morphology have the same spectrum, scarlet models exist in a partially deconvolved frame with the seeing of a well sampled but narrow Gaussian. The initial spectrum of each source is determined using a least squares fit of each monochromatic morphology, convolved with the difference kernel in each band to match the image, for each component. It then uses proximal-ADAM (PADAM; [P. Melchior et al. 2019](#)) to iteratively update the spectrum and morphology with the given constraints until convergence or a maximum number of iterations is reached. It should be noted that although we do use symmetry to initialize the scarlet models, we do not implement a symmetry constraint and the final models are not guaranteed to be symmetric. The models are stored as the `object_model_data` data product, which contains all of the blends for a single patch. Like the single-band deblender, the `scarlet_lite` models are only used as templates to redistribute flux from the image and all measurements are performed on the flux redistributed models.

5.5. Source Measurement

After sources are *detected* (5.3) and optionally *deblended* (5.4), the source measurement tasks are responsible for applying a suite of measurement *plugins* on the deblended pixels for each source. Centroiders, shape measurements, and photometry algorithms are all implemented as measurement plugins.

We also distinguish between measurement on the original detection image (`SingleFrameMeasurementTask`) vs. measurement on a different image from the original detection (`ForcedMeasurementTask`). Measurement could be performed on a single-visit image, a coadd of multiple images, or a difference of images: from the perspective of a measurement plugin, there is no difference between these cases. *Forced measurement* is performed on one image using a “reference” catalog of sources that were detected on another image.

The measurement tasks, plugin base classes, and a suite of standard common plugins are defined in the `meas_base` package, including (but not limited to):

5.5.1. Framework Mechanics

Plugins are enabled or disabled in a measurement task via the task’s configuration, and each plugin has its own configuration nested within the task configuration. When a measurement task is constructed, it constructs instances of its enabled plugins, providing them a schema object that they can use to declare and document their output columns. Each plugin is responsible for defining and filling in columns in the output source catalog, and almost all plugins include columns for uncertainties and at least one flag column to report failures.

Measurement plugins often depend on each other, and must be run in a particular order. Rather than creating a directed acyclic graph to denote the dependencies, the plugins are batched and are run in any order within a batch. The batch order is defined by the `getExecutionOrder` method, with smaller execution numbers being run first. `BasePlugin` defines a list of named constants for particular cases:

1. `CENTROID_ORDER` for plugins that require only footprints and peaks
2. `SHAPE_ORDER` for plugins that require a centroid to have been measured
3. `FLUX_ORDER` for plugins that require both a shape and centroid to have been measured.

The measurement system also provides a *slot* system for predefined aliases to allow a plugin to get a value

without knowing exactly what plugin originally computed that value, e.g., `slot_Centroid` could point to `base_SdssCentroid`, or some other plugin that measures centroids.

While the measurement tasks and plugin interfaces are pure Python, most concrete measurement plugins are implemented in C++, since they need to loop pixels.

When a measurement task is run, it starts by making an empty `SourceCatalog` (from the “`afw.table`” library, see 4) with the plugin-defined schema and one row for each of the `Footprint` objects returned by previous detection and deblending tasks. It then temporarily replaces all pixels within `Footprints` by random noise. As the task loops over each row in the output catalog, that source’s pixels are restored – either to the original `Exposure` pixels for isolated or otherwise un-deblended sources, or to the deblender’s `HeavyFootprint` values for deblended children – and the plugins are called in execution order. Each plugin is given the full modified `Exposure` and a row of the output catalog to fill in. Note that plugins are *not* limited to using only the pixels within a `Footprint`; they get to decide themselves which pixels to use. After each source is measured, the task replaces its pixels with noise again, allowing the next source to be measured independently.

5.5.2. Aperture Corrections

With many different measures of photometry available for both stars and galaxies, establishing a consistent internal photometric system is a challenge, even before we consider the problem (covered in §5.8) of mapping that system to absolute photometry via an external reference catalog.

In addition, standard PSF modeling (§5.7) generally focuses on the core of the true PSF, which is only adequate for photometry for galaxies and fainter stars (while the wings of the PSF matter for bright galaxies as well, the semi-arbitrary definition of the boundary of such galaxies is typically a bigger source of photometry uncertainty).

Our solution to both of these problems is *aperture corrections*, in which we apply multiple photometry algorithms to a suite of isolated bright stars on each detector (by default, the same ones used to build the PSF model), compute the ratio of each algorithm to a standard one (a background-compensated top-hat aperture flux), and then interpolate that ratio using Chebyshev polynomials to other positions on the detector. The measured fluxes of other objects are then scaled by that interpolated ratio. This essentially forces all algorithms to produce the same results (on average) on stars.

It is clear that this approach is not exactly correct for galaxies or other extended sources in general, but we believe it is usually better than not correcting the galaxy photometry at all (it is, after all, the right thing to do for barely-resolved galaxies).

These scheme has two other serious limitations. The first is that many of the measurements we need to correct are noisy, even on bright stars, and the interpolated ratios can add significant uncertainty into the final fluxes. The second is that aperture corrections cannot in general be coadded along with the images (unlike PSF models); they are not a linear function of the data. Coadding the ratios *is* a valid first-order approximation in the limit where the photometry does not depend on the PSF and the PSFs of the contributing images are similar, and this is what we do at present, but we cannot currently quantify the error this approximation introduces. Our understanding is that using much larger PSF models is the cleanest solution to these problems, but the additional degrees of freedom that would entail may be hard to constrain with the information available.

5.5.3. Sky Objects

TODO

5.5.4. Standard Measurement Plugins

TODO: highlight important `meas_base` algorithms.

5.5.5. Gaussian Aperture and PSF Photometry

`meas_extensions_gaap` implements the Gaussian Aperture and PSF photometry (GAaP) algorithm (K. Kuijken 2008). It is an aperture photometry algorithm designed to obtain consistent colors of extended objects (i.e., galaxies). This is done by weighting each (pre-seeing) region of a galaxy by the same pre-defined 2D Gaussian function in all the bands and is thus largely insensitive to the seeing conditions in the different bands. In practice, this is done by first convolving each object by a kernel (using the same tools described in §6.3) so that the PSF is Gaussian and is larger by about 15% (this is configurable). As a second step, each Gaussianized object is then weighted with a Gaussian aperture so that the effective pre-seeing Gaussian aperture is the same for all objects in all the bands. The plugin is configured to use a series of circular Gaussian apertures, an elliptical Gaussian aperture (optionally) that matches the shape of the object in the reference band.

Although the two-step approach is motivated by the original implementation in K. Kuijken (2008), the implementation of this algorithm within the broader context of the measurement framework makes it different from the implementation used in the Kilo-Degree Survey (KiDS; A. H. Wright et al. 2024). In particular, because

neighboring objects are replaced with noise before measurement, Gaussianization of the PSF does not result in increased blending as mentioned in Appendix A2 of K. Kuijken et al. (2015). Furthermore, the uncertainty handling is different. Correlations in noise introduced due to PSF-Gaussianization is included in the uncertainty estimates. However, because only per-pixel noise variance is tracked, the noise treatment is forced to assume that the noise is uncorrelated to begin with which is not true on the coadds. See A. Kannawadi (2022) for more details on the implementation details.

Note that this measurements from this plugin do not produce total fluxes, but should only be used to obtain colors. For total fluxes, measurements from `cModel` or `MultiProFit` (c.f., §5.6.2) are recommended.

5.5.6. Kron Photometry

The `meas_extensions_photometryKron` implements Kron photometry (R. G. Kron 1980). Our Kron implementation uses a scaled version of the HSM shapes (See 5.5.7) to form an elliptical aperture, and then scales to $2.5 \times R_{\text{text}}$, where R_{text} is the first radial moment.

Our implementation does not correct for the PSF in any way; this means its outputs should only be used for very well-resolved galaxies. We do not expect our Kron photometry to be competitive with most of our other galaxy photometry algorithms in robustness or precision, but it may be useful for comparison with external measurements.

5.5.7. HSM Shapes

The `meas_extensions_shapeHSM` package contains the plugins to measure the shapes of objects. The plugins measure the moments of the sources and PSFs with adaptive Gaussian weights. The algorithm was initially described in C. Hirata & U. Seljak (2003) and was modified later in R. Mandelbaum et al. (2005). The implementation of these algorithms lives within the `hsm` module of the `GalSim` package (B. T. P. Rowe et al. 2015). `meas_extensions_shapeHSM` now interacts directly with the Python layer of `GalSim` to make the measurements.

The base plugin for measuring moments is the `HsmMomentsPlugin` and is the parent class of the `HsmSourceMomentsPlugin` and `HsmPsfMomentsPlugin` which are specialized to measure on the sources (and objects) and PSFs respectively. `HsmSourceMomentsRoundPlugin` is a further specialized plugin that measures the moments with circular Gaussian weights instead of the elliptical ones in `HsmSourceMomentsPlugin`. The `HsmPsfMomentsDebiasedPlugin` adds noise to the PSF image to degrade it to have the same signal-to-noise

ratio (SNR) as the source image. This makes the ellipticity calculated from this plugin have the same bias as the source ellipticity. The PSF moments from this plugin should be used when calculating ellipticity residuals so the bias is largely cancelled. Having the various specializations as distinct plugins allows an object to be measured under different configurations simultaneously and included in the output catalogs.

In addition to the plugins that measure (adaptive) weighted moments, there are also a series of `HsmShape` plugins to estimate the PSF-corrected ellipticities of objects. In particular, the outputs from `HsmShapeRegaussPlugin` have been used to measure weak gravitational lensing signals in the Hyper Suprime-Cam SSP data (R. Mandelbaum et al. 2018; X. Li et al. 2022).

In addition to the second moments that characterize the size and ellipticity of the PSF, higher-order moments — those beyond second order — capture more subtle aspects of the PSF shape, such as skewness, kurtosis, and other asymmetric or non-Gaussian features. The `HigherOrderMomentsPSFConfig` is a plugin within `meas_extensions_shapeHSM` to calculate the higher order moments of the PSF models whereas `HigherOrderMomentsSourcePlugin` calculates that of the sources (and objects). The definitions of the higher moments are given in T. Zhang et al. (2023). These moments are measured in normalized coordinates, where the normalized x -axis is along the major axis and the normalized y -axis along the minor. Such a normalization implies that the moments are dependent only on features of the light profile beyond second moments, and does not scale with the flux, position, size or orientation of the object. It is designed to be used in conjunction with the `HsmShapePlugin` and `HsmPsfMomentsPlugin` plugins, which measure the second moments for the normalization, and provides the HSM adaptive Gaussian kernel. By default, we compute the third and fourth order moments of the source and PSF images.

5.6. Trailed Sources

TODO: rewrite this to avoid verbatim-copy and check for accuracy.

Implemented in the `meas_extensions_trailed-Sources` package, this component measures the properties of trailed sources, such as those caused by moving objects like asteroids or satellites. It measures the length, angle, flux, centroid, and end points of a trailed source using the P. Vereš et al. (2012) model. This plugin is designed to refine the measurements of trail length, angle, and end points from `NaivePlugin`, and of flux and centroid from previous measurement algo-

rithms. [P. Vereš et al. \(2012\)](#) derive a model for the flux in a given image pixel by convolving an axisymmetric Gaussian with a line. The model is parameterized by the total flux, trail length, angle from the x-axis, and the centroid. The best estimates are computed using a chi-squared minimization.

5.6.1. CModel Galaxy Fitting

The `meas_modelfit` package’s CModel algorithm is a reimplemention of the SDSS galaxy-model fitting approach, in which we fit a PSF-convolved elliptical exponential profile and de Vaucouleurs profile to each object separately, and then fit a linear combination of the two with the ellipse parameters held fixed. This is not as principled as a true bulge-disk decomposition, in which both models are fit simultaneously, but since each fit has fewer degrees of freedom, it can nevertheless work better on low signal-to-noise or poorly-resolved galaxies, where the degeneracies in a bulge-disk decomposition might otherwise lead to completely unphysical parameters. In fact, in practice, CModel is better thought of as a crude single-Sersic fit than a kind of decomposition.

Because these are PSF-convolved models, we expect CModel to provide decent photometry (including for colors) on both small and well-resolved objects. At present, CModel flux uncertainties are known to be severely underestimated; at least some of this is due to the fact that the ellipse parameter uncertainties cannot easily be propagated into the flux uncertainties in the final fit.

The CModel code is essentially unchanged from the version used in the HSC pipelines; and additional details can be found in [J. Bosch et al. \(2018\)](#).

5.6.2. MultiProFit Galaxy Fitting

MultiProFit is a package for Gaussian mixture model fitting ([D. S. Taranu 2025](#)). MultiProFit is primarily used to provide multiband Sersic model fits to objects using all available coadds. The `multiprofit` package is a standalone Python-only package that provides the interfaces for astronomical object fitting. `multiprofit` depends primarily on `gauss2d_fit`, a standalone C++ package with Python bindings for fast evaluation of Gaussian mixture model likelihoods and gradients thereof. `gauss2d_fit` in turn is an extension of `gauss2d`, providing additional classes for parameters with arbitrary limits and transformations from the `modelfit_parameters` header-only C++ library. All of these packages are included in the science pipelines but can also be installed independently, as `multiprofit` only depends on other standalone packages like `pex_config`.

The `meas_extensions_multiprofit` package contains pipeline tasks (with interfaces defined in

`pipe_tasks`) necessary to run `multiprofit` on coadded and deblended images. The first of these tasks fits a Gaussian mixture model to the PSF model image at the location of each object in a patch. This procedure is similar to the shapelet PSF fitting functionality in `meas_modelfit`. The main differences are that the components are pure Gaussians (shapelet parameters are not supported), can have independent shapes, and are constrained to have integrals summing to unity (i.e., they are normalized). Currently, only a maximum of two components are supported; this limitation may be removed in the future.

The remainder of the tasks in `meas_extensions_multiprofit` use the Gaussian mixture PSF model to fit a PSF-convolved model to all objects in a given patch, for all available bands. Convenient tasks are available for a variety of models, including a single Sersic, as well as multi-component bulge-disk models with an optional central point source component. In all cases, the structural parameters for each component are band-independent, with a separate total flux parameter for each band. That is, individual components do not have intrinsic color gradients (although the convolved models might, if the PSF parameters vary by band).

5.6.3. Reliability Scoring

TODO: move this out of the measurement subsection

The `meas_transiNet` package determines a numerical score for input cutout images using pre-trained machine-learning models. Image differencing may produce false detections, so time-domain surveys characteristically use machine learning classifiers to distinguish astrophysical sources from artifacts (“Real/Bogus;” e.g., [J. S. Bloom et al. 2012](#); [D. A. Goldstein et al. 2015](#); [D. A. Duev et al. 2019](#)).

The `meas_transiNet` defines “model packages” that consist of a python architecture class, a PyTorch ([A. Paszke et al. 2019](#)) weights file, and associated metadata. The inference task may be configured to load a model package from disk or from the Butler.

The `RBTransiNetTask` PipelineTask takes as input three square cutouts of configurable size from the science, template, and difference images centered on the location of a source. These images are concatenated, batched into Torch blobs, and passed to the model for inference. Either CPU or GPU backends may be used for inference. The output of the task is a single float ranging from 0–1 for each cutout triplet, with higher values indicating that the DIASource is more likely to be astrophysical. These reliability scores are then joined with the DIASource catalogs by a later transformation task. Detailed discussion of the model architecture, training,

and performance will be presented in T. Acero Cuellar et. al (in prep.).

5.7. PSF Modeling

Point-spread function (PSF) modeling in the LSST Science Pipelines is largely delegated to external libraries, albeit with considerable wrapper code to adapt them to a common interface. We use both a heavily modified version of PSFEx (E. Bertin 2011) and Piff (M. Jarvis et al. 2021a,b) in our production pipelines. PSFEx is faster, and is used in nightly alert processing and as to obtain a preliminary model in data release processing, while Piff is used in a second, more careful round of PSF estimation in data release processing for better accuracy (especially for weak gravitational lensing).

5.7.1. *meas_extensions_psfex*

The `meas_extensions_psfex` package provides an interface to our own library version of the PSFEx command-line tool. At its core is `PsfexPsfDeterminerTask`, which prepares these selected stars for input into PSFEx, calls into the library, and converts the output into an LSST-specific PSF object (`PsfexPsf`). Key parameters such as spatial interpolation order and oversampling ratio are controlled via `PsfexPsfDeterminerConfig`.

For each CCD in the focal plane, PSFEx independently models the PSF as a linear combination of basis vectors and captures spatial variation using polynomial interpolation. `PsfexStarSelectorTask` offers a built-in mechanism for star selection using strict cuts on signal-to-noise ratio, FWHM range, ellipticity, and quality flags.

5.7.2. *meas_extensions_piff*

The `meas_extensions_piff` package is a wrapper around the Piff package. Piff is a modular package that supports various PSF models, interpolation schemes, and coordinate systems. It can operate on a per-CCD basis or over the full field of view, as indicated by its name. The implementation within `meas_extensions_piff` does not exploit the full modularity of Piff; instead, it closely follows the method used for cosmic shear analysis like in DES (M. Jarvis et al. 2021b; T. Schutt et al. 2025), but it has been designed to make it easy to add support for more options as needed.

The PSF model utilized is a `PixelGrid`, and the interpolation is performed using `BasisPolynomial` interpolation (M. Jarvis et al. 2021b). Modeling is executed per CCD and can employ either pixel or sky coordinates. A

key difference from PSFEx is that Piff implements outlier rejection based on chi-squared criteria (see M. Jarvis et al. 2021b, for more details).

Most of the configuration described here is adjustable through the `PiffPsfDeterminerConfig`. Some important features that were implemented by M. Jarvis et al. (2021b) and T. Schutt et al. (2025) have not yet been enabled but will be available in the near future. M. Jarvis et al. (2021b) configure Piff to fit in sky coordinates with a world-coordinate system transform (WCS) that includes CCD distortions such as tree rings. `meas_extensions_piff` is capable of doing the same, but our models do not yet include CCD distortions, and hence thus far we have not used this approach in our production configuration, as it doesn't significantly improve the quality of the PSF models. Additionally, although T. Schutt et al. (2025) incorporated a color correction to account for chromatic effects on the PSF, this correction has not yet been implemented in `meas_extensions_piff`.

5.8. Astrometric and Photometric Calibration

Astrometric and photometric calibration in data release processing is performed in two very different steps. First, astrometric and photometric solutions for each detector are fit independently to a reference catalog, which is sufficient to enable matching and filtering for more sophisticated later algorithms and validation. In the nightly alert processing, we expect to be able to use a high-density reference catalog produced by the most recent Rubin data release, and these single-detector fits represent the final calibrations. In data release processing, a much more sophisticated final astrometric and photometric transform are then fit to catalogs from multiple epochs at once.

The Rubin pipeline uses the Starlink AST library (D. S. Berry et al. 2016) for persisting, composing, and evaluating coordinate transformations. We have our own `BoundedField` class in the `afw` package for representing (among other things) photometric calibrations, usually with Chebyshev polynomials. These can also be multiplied and even mixed with AST-backed transform objects to represent pixel-area corrections.

5.8.1. *Single-Frame Astrometric Calibration*

Single-frame astrometric fits are performed by `AstrometryTask` in `meas_astrom`. This task matches a catalog of sources detected and measured on an image to a reference catalog and solves for the *World Coordinate System* (WCS) of the image. Matching and WCS fitting are performed iteratively, to reject astrometric outliers. The matcher is either the optimistic (`MatchOptimisticBTask`) or pessimistic (`MatchPessimistic-`

BTask) matcher from V. Tabur (2007), with the pesimistic matcher used by default due to better performance on dense fields; see (C. B. Morrison 2018) for details. The WCS fitter can be a simple affine model on top of fixed camera geometry, as in `FitAffineWcsTask`, or a FITS TAN-SIP WCS (D. L. Shupe et al. 2005), as in `FitTanSipWcsTask` or `FitSipDistortionTask`. We default to fitting the simple affine model because we have good distortion models for most of the instruments we support, often from previous fitting with `gbdes` (§5.8.2), and hence we do not need the extra degrees of freedom provided by a TAN-SIP model.

5.8.2. GBDES

The final astrometric solution is fit by `GbdesAstrometricFitTask` in `drp_tasks`, which runs the `wcsfit` fitter from the `gbdes` package (G. M. Bernstein 2022; G. M. Bernstein et al. 2017) on the ensemble of images in a given band overlapping with a given tract. This task fits a per-detector polynomial distortion model, a per-exposure polynomial distortion model, and position for all the isolated star sources in the component images. This is done by first associating all isolated point sources in the input images and matching them with an external reference catalog. The model is then fit by iterating between fitting the per-detector and per-exposure polynomial models, and recalculating the best-fit solution for the object positions.

The task can be configured to fit either a two-parameter (position on the sky) or five-parameter (position, proper motion, and parallax) solution for the input objects. Correcting for differential chromatic refraction is another configurable option.

There are also options to run variants of the main `GbdesAstrometricFitTask`: `GbdesAstrometricMultibandFitTask` fits images from multiple bands at once, in which case the per-detector distortion model is also per-band; `GbdesGlobalAstrometricFitTask` removes the restriction to a single tract and fits images regardless of their location on the sky by splitting the images into contiguous groups; `GbdesGlobalAstrometricMultibandFitTask` combines these two options.

Lastly, the per-detector polynomial model fit by the task is also used to build a camera distortion model, which can be fed back into single-frame modeling or into the `gbdes` fit for other data. For use in single-frame modeling, the `BuildCameraFromAstrometryTask` subtask is used to build an `afw` Camera object out of the native polynomial model.

A full description of astrometric calibration in the pipeline is given in C. Saunders (2024).

5.8.3. Single-Frame Photometric Calibration

Single frame photometric calibration is performed by `PhotoCalTask`, in the `pipe_tasks` package. The catalog to be calibrated is down-selected to contain only bright ($S/N > 10$), well measured, PSF-like sources which are then matched to a reference catalog. The matched sources have their instrumental fluxes converted into rough magnitudes, which are iteratively compared with the reference catalog magnitudes using a sigma-clipping algorithm, to fit a single magnitude zero point to the whole image. Precomputed color terms can also be applied to the reference catalog fluxes when needed.

5.8.4. FGCM

Global photometric calibration is computed via the Forward Global Calibration Method (FGCM D. L. Burke et al. 2018), as adapted for LSST (P. Fagrellius & E. Rykoff 2025). This global calibration algorithm makes use of repeated observations of stars in all *ugrizy* bands, combining a forward model of the atmospheric parameters with instrumental throughputs measured with auxiliary information. In this way we simultaneously constrain the atmospheric model as well as standardized top-of-atmosphere (TOA) star fluxes over a wide range of star colors, including full chromatic corrections from the instrument and atmosphere.

Running `fgcmcal` first requires generating a look-up table. The input to the look-up table includes the effect of a MODTRAN (A. Berk et al. 1999) atmospheric model at the elevation of the observatory, as well as the throughput as a function of wavelength and position from the optics, filters, and detector quantum efficiency. The quality of the output (in terms of repeatability of bright isolated stars across a wide range of colors) depends on the knowledge of the instrumental throughput.

The primary goal of `fgcmcal` is to provide a uniform relative photometric calibration of the survey. For “absolute” (relative) calibration, a reference catalog can be used as an additional constraint on the fit. Thus, the overall throughput output by `fgcmcal` depends on the reference catalog. This can be checked with (e.g.) specific white dwarfs or CALSPEC (R. C. Bohlin 2007) stars in the survey. However, the relative spatial and chromatic calibration of the `fgcmcal` calibration means that the absolute calibration reduces to a set of 6 numbers (one for each band, or one overall throughput and 5 absolute colors).

5.8.5. jointcal

`jointcal` fits both astrometry and photometry across multiple exposures of large mosaic cameras, fitting for both the true star positions/fluxes, and the distortions

caused by the telescope and instrument. `jointcal` is no longer used by the LSST camera pipeline, but is available for use by cameras that are not supported by `gbdes` and/or `fgcmcal` (for example, DECam). More details on the `jointcal` algorithm are available in J. Parejko & P. Astier (2018).

5.9. Catalog Schemas

Pipeline products must be transformed from the internal data model to the public data model defined in M. Jurić et al. (2023). A set of YAML files in the `pipe_tasks`⁸ repository are used for transforming the internal pipelines representation of the data to a standardized parquet output format. These parquet files are continually validated against an appropriate schema to ensure that the column names and types are correct as the pipelines codebase evolves, as described in Section ?? . For data previews and releases, the parquet files are then ingested into the Qserv database, where the data catalogs are stored.

The public data model is defined by a set of files in the Felis (J. McCormick et al. 2024) YAML format which describe the schema of a data catalog, including its tables, columns, constraints and metadata. These are collectively referred to as the Science Data Model (SDM) schemas. The YAML files are managed with the `sdm_schemas` repository⁹ with all changes validated by GitHub workflows. The schemas corresponding to Science Pipelines output are continually evolving with the pipelines codebase, so, for instance, column names and types may be updated to reflect changes to the internal data model. Schemas for data previews and releases represent a snapshot of the public data model at the time of the release and would typically only be updated with bug fixes, minor changes, or updates and additions to the metadata. For public-facing data catalogs, the Felis representation is used to generate a TAP_SCHEMA database describing the tables and columns available in the TAP service (P. Dowler et al. 2019) and serves as a source of documentation. User data access to the data catalogs is provided primarily through the Astronomical Data Query Language (ADQL) through the TAP service, with the tables and columns being validated against the TAP_SCHEMA as part of the query execution process.

6. HIGH-LEVEL TASKS AND PIPELINES

In this section, we describe how the reusable components of 5 are assembled into pipelines for generating

nightly alerts immediately from just-observed data, producing cumulative data releases roughly every year, and creating calibrations that are input to both on many cadences in between. This also includes some description of algorithmic details that are specific to certain pipelines.

6.1. Single-Frame Processing and Calibration

6.1.1. CalibrateImage

`CalibrateImageTask`, from the `pipe_tasks` package, performs “single frame processing” on a post-ISR (§5.1) single detector exposure. We repair and mask cosmic rays and defects, perform an initial set of detection (§5.3) and measurement (§5.5) passes to estimate the image Point Spread Function (PSF), compute an astrometric (§5.8.1) and photometric (§5.8.3) calibration, and compute summary statistics on the resulting nanojansky-calibrated exposure and catalog. The primary user-facing outputs of this task are the photometrically calibrated, background-subtracted `preliminary_visit_image`, the calibrated `preliminary_visit_image_background` that was subtracted from it, and `single_visit_star_unstandardized`, a catalog of bright point-like sources that were used as inputs to calibration, with only a small number of measurements performed on them. As this task only processes a single detector, it is used by both DRP (§6.7) and AP (§6.6), though with different configurations (AP is focused on latency, while DRP performs more measurements). For more details of the exact steps performed in this task, see the `CalibrateImageTask` pipelines documentation.

6.1.2. ReprocessVisitImage

In DRP, `ReprocessVisitImageTask`, from the `drp_tasks` package, takes the outputs of the global astrometric and photometric models, a visit-level background, and PSF model and re-runs detection and measurement on the post-ISR single-detector exposure. The primary user-facing outputs of this task are the photometrically calibrated, background-subtracted `visit_image`, the calibrated `visit_image_background` that was subtracted from it, and `source_unstandardized`, a catalog of all sources detected to 5-sigma, with all relevant measurements performed on them (this is standardized and consolidated into the `source` per-visit catalog). Besides its primary purpose of performing detection and measurement using the “best” available inputs, this task also allows us to only have to save a small number of relatively small-sized intermediate products in order to regenerate the `visit_image` from a `raw` image, reducing long-term storage needs. For more details of the exact

⁸ https://github.com/lsst/pipe_tasks

⁹ https://github.com/lsst/sdm_schemas

steps performed in this task, see the [ReprocessVisitImage pipelines documentation](#).

6.2. Coaddition and Object Tables

6.2.1. Coaddition Tasks

Coadded images are used as static-sky templates for image subtraction and detecting and measuring faint sources. The coaddition process is divided into two main stages: resampling the input images onto a common projection and stacking those resampled images into a single coadd. Each stage is implemented via configurable tasks that allow the pipelines to be adapted for different instruments and observing strategies. The first step in coaddition is to resample each single-epoch exposure onto a common projection defined by a `skymap` (see section 4), one `patch` at a time. This step is performed by the following **Tasks**:

- **MakeDirectWarpTask** performs a straightforward resampling of calibrated exposures.
- **MakePsfMatchedWarpTask** also convolves them to a configurable, common model-PSF. This PSF-homogenized variant is useful when the scientific goals require uniform PSF properties across the coadd and is used for artifact rejection during coaddition.

All warping tasks use a configurable interpolation kernel. A 5th-order Lanczos kernel is used by default, balancing fidelity and computational efficiency, with a nearest-neighbor kernel for the integer bitmask. Input images are geometrically transformed using the World Coordinate System (WCS) and interpolated onto the target projection defined by a tract and patch geometry. Each resulting resampled image is called a **warp**.

Once warps are generated, they are stacked into a final coadd by **AssembleCoaddTask** or one of its subclasses. The default implementation, **CompareWarpAssembleCoaddTask**, performs outlier rejection to remove transient artifacts such as cosmic rays, ghosts, satellite trails, and moving objects. The algorithm compares pixel values across epochs and masks those that significantly deviate from the expected distribution. The artifact rejection algorithm is detailed in [Y. AlSayyad \(2019\)](#). By default, weights for stacking are derived from the inverse of the average variance of each warp, with optional filters on PSF quality and seeing. The stacked image is accompanied by a mask plane and variance map, and the set of input PSF models is combined into a spatially-varying coadd PSF model (**CoaddPsf**) to serve as the PSF model for the coadd.

6.2.2. Object Catalog Production

TODO

6.3. Difference Image Analysis

Image subtraction for transient/variable detection and analysis is implemented in the `ip_diffim` package, and is divided into three steps. While the image subtraction system could be used to work on other combinations, we primarily focus on subtracting a “template” coadd from a single-visit “science” image. First, the template coadd is warped by **GetTemplateTask** to the WCS and bounding box of the science image. Then the warped template is subtracted from the science image using one of several available algorithms in **SubtractImagesTask**, which produces a temporary difference image. Finally, peaks are detected on the difference image and **DIASources** (“difference image analysis sources”) are measured in **DetectAndMeasureTask**. The final difference image with updated mask planes is written along with the **DIASource** catalog.

6.3.1. PSF Matching and Subtraction

The primary implementation of image subtraction used by **SubtractImagesTask** is based on [C. Alard & R. H. Lupton \(1998\)](#), and uses spatially-varying Gauss-Hermite basis functions for the fit. The PSF-matching kernel can be constructed for either the science or the template image, and the resulting difference image is decorrelated [D. J. Reiss & R. H. Lupton \(2016\)](#). Optionally, the science image can be preconvolved with its own PSF before PSF-matching, producing a Score image analogous to [B. Zackay et al. \(2016\)](#).

6.3.2. DIA Detection and Measurement

Positive and negative peaks are detected by thresholding the Score image if it is available. Otherwise, the difference image is smoothed with a Gaussian of the same width as the PSF of the science image, and thresholds are taken on the smoothed image. Contiguous pixels around each peak that are statistically brighter than the background are grouped into source footprints, and any overlapping footprints are merged. Footprints that contain both a positive and a negative peak are fit as dipoles. The dipole fit simultaneously solves for the negative and positive lobe centroids and fluxes using non-linear least squares minimization. **DIASources** that are not classified as dipoles instead fall back on an SDSS-style centroid ([J. R. Pier et al. 2003](#)). Finally, all configured measurement plugins are run, including HSM shape measurements (Section 5.5.7) and a trailed-source fit.

6.3.3. Filtering Non-astrophysical DIASources

A subset of the DIASources detected during `detectAndMeasureDiaSource` are expected to be non-astrophysical. These may be due to a variety of causes, including uncorrected instrument signatures, diffraction spikes and wings of bright stars, unmasked cosmic rays, algorithmic failures in background subtraction or image differencing, and satellites or space debris orbiting the Earth. The latter can leave long streaks that cross one or more LSSTCam detectors (J. A. Tyson et al. 2020; I. Hasan et al. 2022).

To reduce the number of DiaSources originating from non-astrophysical sources, we apply a multi-prong approach:

- During detection, we remove sources that have pixel flags characteristic of artifacts and unusable measurements, such as when all of the central pixels of a DIASource are saturated.
- In AP, we identify and filter out DIASources that are spatially and temporally coincident with the predicted positions of catalogued artificial satellites.
- DIASources with negative forced fluxes on the direct science image characteristic of bad background subtraction are filtered. This and subsequent catalog filtering is carried out by `ap_association.FilterDiaSourceCatalogTask`.
- Trailed DIASources consistent with a rate of motion greater than 10 degrees per day are filtered.
- Blank sky sources used for noise estimation are removed from the output catalog.

No pixels are altered or redacted during the steps above. The machine-learned reliability score (§5.6.3) provides a further diagnostic that users can apply to remove remaining artifacts. These approaches are intended to enable science users to control the tradeoff between completeness and purity during their analysis while ensuring that the pipelines and databases can maintain their required performance.

6.3.4. Source Association

The `ap_association` package contains multiple tasks for standardizing newly detected DIASources and associating them into “DIAObjects”. Standardization converts the output catalogs from difference imaging to the format specified in `sdm_schemas` (§5.9), and applies filtering consistent with (W. O’Mullane et al. 2024). Once

DIASource catalogs are standardized, they are associated to DIAObjects in either of two modes: Data Release Production (DRP) or Alert Production (AP). Both implementations use the Pessimistic Pattern Matcher B (C. B. Morrison 2018) to score and match DIASources, but differ in how DIAObjects are stored and how visits are ordered.

- **DRP** association loads all DIASource catalogs from a set time period overlapping a single patch at once, and creates new DIAObjects for matched DIASources from all visits simultaneously.
- **AP** association processes a single visit at a time, and creates new DIAObjects incrementally from unassociated DIASources. DIAObjects and their associated DIASources are stored in the Alert Production Database (APDB; §6.3.5).

After association, an additional filtering step may be applied to DIASources with no matched DIAObject or Solar System object (§6.4). Properties of the source such as its reliability score (§5.6.3, source flags, or signal-to-noise cuts may be used to drop detections that are likely to be false detections and avoid creating erroneous new DIAObjects.

6.3.5. Alert Production Database (APDB)

The Alert Production Database (APDB; A. Salnikov & J. McCormick 2024)) supports SQL, Postgres, and Cassandra database formats. The previous history of DIAObjects, DIASources, and DiaForcedSources for the region containing the science image is loaded with `LoadDiaCatalogsTask`, which are passed to `DiaPipelineTask` for association. Loading is split from the association step to enable preloading of catalogs from the database in Prompt Processing during the interval when the next visit has been scheduled but the images have not yet been taken. When AP-style association is run outside of Prompt Processing, it is therefore essential to process all association tasks in strict visit order to prevent loading catalogs from the APDB prematurely and losing DiaObject history in association.

6.3.6. Alert Generation

In order to enable real-time science, the AP pipelines generate alert packets for each detected DIASource. These packets are serialized in Apache Avro¹⁰ format and then transmitted to community alert brokers via Kafka for further processing. M. Patterson et al. (2020) provides a high-level overview of the alert system.

¹⁰ <https://avro.apache.org/>




Figure 6. Detection, attribution, linking, submission and precovery of moving sources within the nightly data: The attribution is performed in real-time by the AP pipelines querying the `mpsky` service with resulting information attached to the alerts and queued for submission to the MPC. The linking is performed in daytime using `helioLinX`, with resulting links queued for submission to the MPC. Fetching of data from the MPC is performed automatically using PostgreSQL replication, with new data triggering recomputation of physical properties and precovery runs in the Daily Data Products Pipeline. Any observations discovered by the precovery procedure are queued for submission to the MPC, using the submission manager tool. The ephemeris cache is precomputed at dusk using `SorchA` and `mpsky`, to enable fast attribution at nighttime. All timings denote design goals.

Within the pipelines, alert packets are constructed by `PackageAlertsTask` within `ap_association`. Alert packets contain the triggering `DIASource` record; the associated `DIAObject` or `SSObject` record; up to twelve months of past history from `DIASources`, `DIAForcedSources`, and/or upper limits; and cutout images of the science, template, and difference images centered at the position of the cutout. Cutouts are provided as FITS images serialized by the astropy `CCDDData` class, and include image, variance, and mask planes along with WCS information and an image of the approximate PSF.

Avro schemas are stored in the `alert_packet` package. They are derived from the corresponding AP schemas in `sdm_schemas` used to instantiate the AP databases.

6.4. Solar System Pipelines

The Solar System Pipeline (SSP; Figure 6) suite is responsible for (i) discovering previously unknown solar system objects by linking together observations (usually `DIASources`) unattributable to static (non-moving) sources, (ii) reporting these to the Minor Planet Center (MPC), (iii) computing basic physical characteristics such as absolute magnitudes and slope parameters for all asteroids where sufficient data is available, and (iv) using the orbits received from the MPC to associate their apparitions in the `DIASource` tables (both in real-time and as precovery).

The core element of the SSP is the linking pipeline, named `helioLinX` (A. Heinze et al. 2023). This code, run in daytime, clusters newly detected `DIAObjects` to search for candidate asteroids. The high-level procedure is to link `DIASource` detections within a night (when on-sky motion is approximately linear) into *tracklets*, to link these tracklets across multiple nights (into tracks) and to fit the tracks with an orbital model to identify those tracks that are consistent with an asteroid orbit. The Rubin implementation of this software (Heinze et al., in prep.) is based on the `HelioLinC` algorithm (M. J. Holman et al. 2018), with the key change being that the clustering is performed not on the sky, but in 3D space. It is designed to be capable of detecting 95% of all Solar System objects whose tracklets are observed over three nights within a 15-night window.¹¹ `helioLinX` is written in C++, but provides a Python API including a `Task API`.

Candidate discoveries with high degree of certainty, as well as re-observations of already known objects, are reported to the Minor Planet Center (MPC) using the observation submission pipeline. The astrometric and photometric data are converted to the PSV variant of the Astrometric Data Exchange Standard (ADES; S. R. Chesley et al. 2017), and submitted via a HTTPS POST API provided by the MPC.

Following processing and validation of newly reported candidates, they're added to the MPC's central database. This database, including the table of orbits as well as observations, is replicated using PostgreSQL logical replication. Following the replication, the Daily Data Products Pipeline recomputes the absolute magnitudes of objects in the `SSObject` table, as well as some auxiliary per-observation information for individual observations (the `SSSource` table).

The replicated orbits and computed absolute magnitudes are utilized to predict positions (ephemerides) and magnitudes of solar system objects in subsequent night. To enable speedy retrieval (on order of 100msec or less) of all objects in a visit, we precompute on-sky locations of all solar system objects, fit Chebyshev polynomials, and build an efficient `HEALpix`-based index allowing for fast lookup. These ephemerides are then served to association pipelines described in Section ?? . This element of the pipeline is based on `SorchA` (computation; S. R. Merritt et al. 2025) and `mpsky` (fast lookup and serving; M. Juric 2014). While still being constructed, a similar service is planned

¹¹ Detailed criteria are specified in the LSST Observatory System Specification (OSS) document `OSS-REQ-0159` (C. F. Claver & The LSST Systems Engineering Integrated Project Team 2018)

for “precovery” – the association of originally missed observations of solar system objects observed earlier in the survey.

Taken together, this suite of pipelines enables Rubin to identify sources consistent with being observations of objects in the solar system (both new and previously known), and makes these data public by reporting their discoveries to the Minor Planet Center and making them available to Rubin users within via the PPDB.

6.5. Calibration pipelines

The high-level pipelines to build calibration products (`cp`) for the LSST cameras are defined in `cp_pipe`¹². They set `IsrTaskLSST` (see Section 5.1) configuration parameters needed for each calibration product, by enabling all the sequential steps of the ISR task up to the step before the correction being generated. In some cases, configurations also specify whether to combine exposures (for bias or dark exposures for instance) and to bin exposures to support diagnostic displays.

Once calibration products are produced, they are “verified” (see C. Waters & E. Rykoff (2025) for more details) using `cp_verify`¹³ pipelines by checking they pass metrics defined in R. Lupton et al. (2025). In this case, verify configuration parameters enable all corrections in the ISR task up to and including the application of the correction being verified. As a result, the calibration products can then be certified to be available in the `butler` and used to ISR an exposure.

6.6. `ap_pipe`

The `ap_pipe` package defines the pipeline(s) to be used for real-time Alert Production processing (K.-T. Lim 2022). These pipelines include instrument signature removal (§5.1), calibration (§5.8), measurement plugins (§5.5), image differencing (§6.3), source association (§6.3.4), and alert generation (§6.3.6). Some of these tasks are shared with the pipelines in `drp_pipe`, but are configured to prioritize speed over strict quality; for example, they use a minimal set of measurement plugins.

`ap_pipe` currently has pipeline variants for LSSTCam, LSSTComCam, LATISS, the Rubin Observatory simulators, Hyper-SuprimeCam, and the Dark Energy Camera. Because these variants serve as testbeds for AP-specific algorithms and configuration settings, they are, as much as possible, the “same” pipeline, differing al-

most entirely in loading instrument defaults from `obs` packages (§3.5). The only other customization is an extra task for handling DECam’s inter-chip crosstalk, which does not have an equivalent for Rubin instruments.

6.7. `drp_pipe`

The `drp_pipe` package defines the pipeline(s) to be used for the annual Data Release Production processing. These pipelines include instrument signature removal (§5.1), calibration (§5.8), measurement plugins (§5.5), coaddition and coadd-processing (§??), image differencing (§6.3), source association (§6.3.4), and global calibration (§5.8). Some of these tasks are shared with the pipelines in `ap_pipe`, but are configured to prioritize accuracy over speed.

`drp_pipe` currently has pipeline variants for LSSTCam, LSSTComCam, LATISS, the Rubin Observatory simulators, Hyper-SuprimeCam, and the Dark Energy Camera. Because these variants serve as testbeds for DRP-specific algorithms and configuration settings, they are, as much as possible, the “same” pipeline, differing primarily entirely in loading instrument defaults from `obs` packages (§3.5). The pipelines for some instruments also disable certain multi-visit calibration routines like FGCM (§5.8.4) and GBDES (§5.8.2), as those can take significant effort to commission on a new instrument, and this has not been done in all cases. The only other customization is an extra task for handling DECam’s inter-chip crosstalk, which does not have an equivalent for Rubin instruments.

7. ANALYSIS TOOLING

7.1. Display Abstractions

The `afw.display` subpackage defines a simple image display abstraction layer that be used to show our image objects via multiple applications and libraries. This includes programmatic control of stretch levels and colors, WCS mapping, semi-transparent bitmask overlays, and simple geometric region support.

There are currently implementations for `matplotlib` (J. D. Hunter 2007), `Firefly` (W. Roby et al. 2020), `SAOImage DS9` (W. A. Joye & E. Mandel 2003), and `Ginga` (E. Jeschke et al. 2013, via `Astrowidgets`). While most of these tools have considerable functionality beyond what our abstraction layer provides, the ability to interact easily with them in a consistent, programmatic way in many different contexts (e.g. DS9 on personal machines, `matplotlib` and `Ginga` in Jupyter notebooks, and `Firefly` in the Rubin Science Platform) is invaluable, and usually the `afw.display` provides all we need.

¹² https://github.com/lsst/cp_pipe and see documentation at <https://pipelines.lsst.io/modules/lsst.cp.pipe/constructing-calibrations.html>

¹³ https://github.com/lsst/cp_verify

7.2. Analysis Tools

The `analysis_tools` package provides a framework to allow reproducible, automatic creation of plots and metrics through a set of configurable, reusable tools that can be used in pipeline execution and interactive analysis. The package allows metrics and plots to be consistently created at various points in the pipeline and ensures that the metrics dispatched to the monitoring dashboard are generated in sync with the archived plots. An example plot, made with HSC data, is shown in Fig. 7. The package was designed to handle the large data volumes and memory requirements that the survey will generate to ensure that the initial QA products required are rapidly made and readily available for fast action on any emergent data quality issues. The individual tools run in the pipelines to calculate the metrics can then be reused in an interactive environment, such as a script or notebook, allowing further investigation into arising issues to reproduce exactly what was originally run.

Further information and examples can be found in [S. L. Reed \(2025\)](#).




Figure 7. An example figure produced as part of the standard processing by analysis tools, the plot is information dense as it is designed for an audience familiar with the outputs but a simplified version can also be produced for talks and publications by setting a config option. The metrics shown in the bottom right of the plot are also saved separately to be displayed by various pieces of QA tooling.

7.3. Source Injection

The `source_injection` package contains tools designed to assist in the injection of synthetic sources into scientific imaging. Source injection is a powerful tool for testing the algorithmic performance of the LSST Science Pipelines, generating measurements on synthetic sources where the truth is known and facilitating subsequent quality assurance checks. Synthetic source generation



Figure 8. An HSC i-band cutout from tract 9813, patch 42, showing before (top) and after (bottom) the injection of a series of synthetic Sérsic sources. Images are 100 arcseconds on the short axis, log scaled across the central 99.5% flux range, and smoothed with a Gaussian kernel of FWHM 3 pixels.

and injection capability is provided by the GALSIM software package ([B. T. P. Rowe et al. 2015](#)). An example showcasing the injection of a series of synthetic Sérsic sources into an HSC i-band image is shown in Figure 8.

Synthetic sources can be injected into any imaging data product output by the LSST Science Pipelines, including visit-level exposure-type or visit-type datasets (i.e., datasets with the dimension `exposure` or `visit`), or into a coadd-level coadded dataset. These injection tasks are defined in `ExposureInjectTask`, `VisitInjectTask` and `CoaddInjectTask`, respectively. Each task operates similarly: read in an injection catalog containing the parameters of the sources to be injected, generate sources using GALSIM, and inject them into the input image. An additional mask plane (`INJECTED` by default) is appended to the image mask to identify pixels which have been touched by injected sources. Optional modifications to the noise profiles of injected sources and the variance plane of the image can also be performed.

With GALSIM we have the capacity to generate synthetic sources of varying profile types, including Gaus-

sian, exponential and Sérsic profiles (J. L. Sérsic 1963, 1968), each convolved with the local PSF. We also have the option to inject scaled versions of the PSF model itself in order to simulate stars. If preferred, a pre-generated FITS image of a source can be injected instead of a model generated by GALSIM, allowing for the injection of complex sources or postage stamp cutouts of real data.

Alongside the primary injection tasks, a suite of helper tools are also provided to optionally assist in the generation of synthetic source catalogs and injection pipelines. Fully qualified source injection pipeline definition YAML files are normally constructed using an existing pipeline as a baseline reference. A user specifies which dataset type they would like to inject synthetic sources into, and the `source_injection` package generates a new pipeline definition YAML file that includes the correctly configured source injection task. By default, all tasks in the pipeline downstream of the point at which source injection occurs are modified such that their connection names are prefixed with `injected_`. This ensures that an injected dataset is not confused with the original dataset when stored together in a common collection.

Once source injection has completed, the source injection task will output two dataset types: an injected image, and an associated injected catalog. The injected image is a copy of the original image with the injected sources added. The injected catalog is a catalog of the injected sources, with the same schema as the original catalog and additional columns describing per-source source injection success outcomes.

8. CONCLUSIONS

The LSST Science Pipelines Software has been developed over 20 years to support the processing of the

Legacy Survey of Space and Time. It has been used to process formal data releases from both Hyper Suprime-Cam and the Rubin Observatory’s LSSTComCam and is now being used to process LSSTCam commissioning data. The software is designed to be extensible and reusable, supporting a plugin architecture that allows new algorithms to be added without modifying the core codebase and includes a dataset tracking system and graph builder that supports scaling of processing on large batch systems.

ACKNOWLEDGMENTS

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

Facilities: Rubin:Simonyi (LSSTCam), Rubin:1.2m (LATISS)

Software: ndarray (<https://github.com/ndarray/ndarray>), astropy (Astropy Collaboration et al. 2022), pytest (H. Krekel 2017), matplotlib (J. D. Hunter 2007), galsim (B. T. P. Rowe et al. 2015), numpy (C. R. Harris et al. 2020), gbdes (G. M. Bernstein 2022), Starlink’s (D. Berry et al. 2022) AST (D. S. Berry et al. 2016), fgcm (<https://github.com/erykoff/fgcm>),

REFERENCES

- Aihara, H., AIsayyad, Y., Ando, M., et al. 2019, Second data release of the Hyper Suprime-Cam Subaru Strategic Program, PASJ, 71, 114, doi: [10.1093/pasj/psz103](https://doi.org/10.1093/pasj/psz103)
- Alard, C., & Lupton, R. H. 1998, A Method for Optimal Image Subtraction, ApJ, 503, 325, doi: [10.1086/305984](https://doi.org/10.1086/305984)
- AIsayyad, Y. 2019, Coaddition Artifact Rejection and CompareWarp, Data Management Technical Note DMTN-080, NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2583441](https://doi.org/10.71929/rubin/2583441)
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package, AJ, 156, 123, doi: [10.3847/1538-3881/aabc4f](https://doi.org/10.3847/1538-3881/aabc4f)
- Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, The Astropy Project: Sustaining and Growing a Community-oriented Open-source Project and the Latest Major Release (v5.0) of the Core Package, ApJ, 935, 167, doi: [10.3847/1538-4357/ac7c74](https://doi.org/10.3847/1538-4357/ac7c74)

- Axelrod, T., Connolly, A., Ivezić, Z., et al. 2004, The LSST Data Processing Pipeline, in American Astronomical Society Meeting Abstracts, Vol. 205, American Astronomical Society Meeting Abstracts, 108.11
- Axelrod, T., Kantor, J., Lupton, R. H., & Pierfederici, F. 2010, An open source application framework for astronomical imaging pipelines, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7740, Software and Cyberinfrastructure for Astronomy, ed. N. M. Radziwill & A. Bridger, 774015, doi: [10.1117/12.857297](https://doi.org/10.1117/12.857297)
- Babuji, Y., Woodard, A., Li, Z., et al. 2019, Parsl: Pervasive Parallel Programming in Python, in Proceedings of the 28th International Symposium on High-Performance Parallel and Distributed Computing, HPDC '19 (New York, NY, USA: Association for Computing Machinery), 25–36, doi: [10.1145/3307681.3325400](https://doi.org/10.1145/3307681.3325400)
- Berk, A., Anderson, G. P., Bernstein, L. S., et al. 1999, MODTRAN4 radiative transfer modeling for atmospheric correction, 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, International Society for Optics and Photonics (SPIE), 348 – 353, doi: [10.1117/12.366388](https://doi.org/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, Astrometric Calibration and Performance of the Dark Energy Camera, PASP, 129, 074503, doi: [10.1088/1538-3873/aa6c55](https://doi.org/10.1088/1538-3873/aa6c55)
- Berry, D., Graves, S., Bell, G. S., et al. 2022, Starlink - The Original and Best, in Astronomical Society of the Pacific Conference Series, Vol. 532, Astronomical Data Analysis Software and Systems XXX, ed. J. E. Ruiz, F. Pierfederici, & P. Teuben, 559
- Berry, D. S., Warren-Smith, R. F., & Jenness, T. 2016, AST: A library for modelling and manipulating coordinate systems, Astronomy and Computing, 15, 33, doi: [10.1016/j.ascom.2016.02.003](https://doi.org/10.1016/j.ascom.2016.02.003)
- Bertin, E. 2011, Automated Morphometry with SExtractor and PSFEx, 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
- Bertin, E., & Arnouts, S. 1996, SExtractor: Software for source extraction., A&AS, 117, 393, doi: [10.1051/aas:1996164](https://doi.org/10.1051/aas:1996164)
- Bloom, J. S., Richards, J. W., Nugent, P. E., et al. 2012, Automating Discovery and Classification of Transients and Variable Stars in the Synoptic Survey Era, PASP, 124, 1175, doi: [10.1086/668468](https://doi.org/10.1086/668468)
- Bohlin, R. C. 2007, HST Stellar Standards with 1% Accuracy in Absolute Flux, in Astronomical Society of the Pacific Conference Series, Vol. 364, The Future of Photometric, Spectrophotometric and Polarimetric Standardization, ed. C. Sterken, 315, doi: [10.48550/arXiv.astro-ph/0608715](https://doi.org/10.48550/arXiv.astro-ph/0608715)
- Bosch, J., Armstrong, R., Bickerton, S., et al. 2018, The Hyper Suprime-Cam software pipeline, PASJ, 70, S5, doi: [10.1093/pasj/psx080](https://doi.org/10.1093/pasj/psx080)
- Boulade, O., Charlot, X., Abbon, P., et al. 2003, MegaCam: the new Canada-France-Hawaii Telescope wide-field imaging camera, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 4841, Instrument Design and Performance for Optical/Infrared Ground-based Telescopes, ed. M. Iye & A. F. M. Moorwood, 72–81, doi: [10.1117/12.459890](https://doi.org/10.1117/12.459890)
- Broughton, A., Utsumi, Y., Plazas Malagón, A. A., et al. 2024, Mitigation of the Brighter-fatter Effect in the LSST Camera, PASP, 136, 045003, doi: [10.1088/1538-3873/ad3aa2](https://doi.org/10.1088/1538-3873/ad3aa2)
- Burke, D. L., Rykoff, E. S., Allam, S., et al. 2018, Forward Global Photometric Calibration of the Dark Energy Survey, AJ, 155, 41, doi: [10.3847/1538-3881/aa9f22](https://doi.org/10.3847/1538-3881/aa9f22)
- Cai, M., Xu, Z., Fan, L., et al. 2025, The 2.5-meter Wide Field Survey Telescope Real-time Data Processing Pipeline I: From raw data to alert distribution, arXiv e-prints, arXiv:2501.15018, doi: [10.48550/arXiv.2501.15018](https://doi.org/10.48550/arXiv.2501.15018)
- Chesley, S. R., Hockney, G. M., & Holman, M. J. 2017, Introducing ADES: A New IAU Astrometry Data Exchange Standard, in AAS/Division for Planetary Sciences Meeting Abstracts, Vol. 49, AAS/Division for Planetary Sciences Meeting Abstracts #49, 112.14
- Claver, C. F., & The LSST Systems Engineering Integrated Project Team. 2018, Observatory System Specifications (OSS), Systems Engineering Controlled Document LSE-30, NSF-DOE Vera C. Rubin Observatory. <https://ls.st/LSE-30>
- DePoy, D. L., Abbott, T., Annis, J., et al. 2008, The Dark Energy Camera (DECam), in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7014, Ground-based and Airborne Instrumentation for Astronomy II, ed. I. S. McLean & M. M. Casali, 70140E, doi: [10.1117/12.789466](https://doi.org/10.1117/12.789466)

- Dowler, P., Rixon, G., Tody, D., & Demleitner, M. 2019, Table Access Protocol Version 1.1., IVOA Recommendation 27 September 2019
doi: [10.5479/ADS/bib/2019ivoa.spec.0927D](https://doi.org/10.5479/ADS/bib/2019ivoa.spec.0927D)
- Duev, D. A., Mahabal, A., Masci, F. J., et al. 2019, Real-bogus classification for the Zwicky Transient Facility using deep learning, *MNRAS*, 489, 3582, doi: [10.1093/mnras/stz2357](https://doi.org/10.1093/mnras/stz2357)
- Esteves, J. H., Utsumi, Y., Snyder, A., et al. 2023, Photometry, Centroid and Point-spread Function Measurements in the LSST Camera Focal Plane Using Artificial Stars, *PASP*, 135, 115003, doi: [10.1088/1538-3873/ad0a73](https://doi.org/10.1088/1538-3873/ad0a73)
- Fagrelus, P., & Rykoff, E. 2025, Rubin Baseline Calibration Plan, Commissioning Technical Note SITCOMTN-086, NSF-DOE Vera C. Rubin Observatory.
<https://sitcomtn-086.lsst.io/>
- Fausti, A. 2023, Sasquatch: beyond the EFD, SQuARE Technical Note SQR-068, NSF-DOE Vera C. Rubin Observatory. <https://sqr-068.lsst.io/>
- Fausti Neto, A., Economou, F., Reuter, M. A., et al. 2024, Sasquatch: Rubin Observatory metrics and telemetry service, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 13101, Software and Cyberinfrastructure for Astronomy VIII, ed. J. Ibsen & G. Chiozzi, 131011M, doi: [10.1117/12.3019081](https://doi.org/10.1117/12.3019081)
- Fernandes, L. A. F., & Oliveira, M. M. 2008, Real-time line detection through an improved Hough transform voting scheme, *Pattern Recognition*, 41, 299, doi: [10.1016/j.patcog.2007.04.003](https://doi.org/10.1016/j.patcog.2007.04.003)
- Flaugher, B., Diehl, H. T., Honscheid, K., et al. 2015, The Dark Energy Camera, *The Astronomical Journal*, 150, 150, doi: [10.1088/0004-6256/150/5/150](https://doi.org/10.1088/0004-6256/150/5/150)
- Goldstein, D. A., D'Andrea, C. B., Fischer, J. A., et al. 2015, Automated Transient Identification in the Dark Energy Survey, *AJ*, 150, 82, doi: [10.1088/0004-6256/150/3/82](https://doi.org/10.1088/0004-6256/150/3/82)
- Gower, M., Kowalik, M., Lust, N. B., Bosch, J. F., & Jenness, T. 2022, Adding Workflow Management Flexibility to LSST Pipelines Execution, arXiv e-prints, arXiv:2211.15795, doi: [10.48550/arXiv.2211.15795](https://doi.org/10.48550/arXiv.2211.15795)
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Array programming with NumPy, *Nature*, 585, 357, doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)
- Hasan, I., Tyson, J. A., Saunders, C., & Xin, B. 2022, Mitigating satellite trails: A study of residual light after masking, *Astronomy and Computing*, 39, 100584, doi: [10.1016/j.ascom.2022.100584](https://doi.org/10.1016/j.ascom.2022.100584)
- Heinze, A., Juric, M., & Kurlander, J. 2023, helioline: Open Source Solar System Discovery Software,
<https://github.com/helioline/helioline>
- Hirata, C., & Seljak, U. 2003, Shear calibration biases in weak-lensing surveys, *MNRAS*, 343, 459, doi: [10.1046/j.1365-8711.2003.06683.x](https://doi.org/10.1046/j.1365-8711.2003.06683.x)
- Holman, M. J., Payne, M. J., Blankley, P., Janssen, R., & Kuindersma, S. 2018, HeliLinC: A Novel Approach to the Minor Planet Linking Problem, *AJ*, 156, 135, doi: [10.3847/1538-3881/aad69a](https://doi.org/10.3847/1538-3881/aad69a)
- HTCondor Team. 2024, HTCondor, 23.0.15 Zenodo, doi: [10.5281/zenodo.14238973](https://doi.org/10.5281/zenodo.14238973)
- Hunter, J. D. 2007, Matplotlib: A 2D Graphics Environment, *Computing in Science and Engineering*, 9, 90, doi: [10.1109/MCSE.2007.55](https://doi.org/10.1109/MCSE.2007.55)
- Ingraham, P., Clements, A. W., Ribeiro, T., et al. 2020, Vera C. Rubin Observatory auxiliary telescope commissioning as a control system pathfinder, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 11452, Software and Cyberinfrastructure for Astronomy VI, ed. J. C. Guzman & J. Ibsen, 114520U, doi: [10.1117/12.2561112](https://doi.org/10.1117/12.2561112)
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, LSST: From Science Drivers to Reference Design and Anticipated Data Products, *ApJ*, 873, 111, doi: [10.3847/1538-4357/ab042c](https://doi.org/10.3847/1538-4357/ab042c)
- Jarvis, M., Meyers, J., Leget, P.-F., & Davis, C. 2021a, Piff: PSFs In the Full FOV,, *Astrophysics Source Code Library*, record ascl:2102.024
- Jarvis, M., Bernstein, G. M., Amon, A., et al. 2021b, Dark Energy Survey year 3 results: point spread function modelling, *MNRAS*, 501, 1282, doi: [10.1093/mnras/staa3679](https://doi.org/10.1093/mnras/staa3679)
- Jenness, T. 2020, Modern Python at the Large Synoptic Survey Telescope, in Astronomical Society of the Pacific Conference Series, Vol. 522, Astronomical Data Analysis Software and Systems XXVII, ed. P. Ballester, J. Ibsen, M. Solar, & K. Shortridge, 541, doi: [10.48550/arXiv.1712.00461](https://doi.org/10.48550/arXiv.1712.00461)
- Jenness, T., Bosch, J., Owen, R., et al. 2016, Investigating interoperability of the LSST data management software stack with Astropy, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9913, Software and Cyberinfrastructure for Astronomy IV, ed. G. Chiozzi & J. C. Guzman, 99130G, doi: [10.1117/12.2231313](https://doi.org/10.1117/12.2231313)

- Jenness, T., Economou, F., Findeisen, K., et al. 2018, LSST data management software development practices and tools, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10707, Software and Cyberinfrastructure for Astronomy V, ed. J. C. Guzman & J. Ibsen, 1070709, doi: [10.1117/12.2312157](https://doi.org/10.1117/12.2312157)
- Jenness, T., Bosch, J. F., Salnikov, A., et al. 2022, The Vera C. Rubin Observatory Data Butler and pipeline execution system, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12189, Software and Cyberinfrastructure for Astronomy VII, 1218911, doi: [10.1117/12.2629569](https://doi.org/10.1117/12.2629569)
- Jeschke, E., Inagaki, T., & Kackley, R. 2013, Introducing the Ginga FITS Viewer and Toolkit, in Astronomical Society of the Pacific Conference Series, Vol. 475, Astronomical Data Analysis Software and Systems XXII, ed. D. N. Friedel, 319
- Joye, W. A., & Mandel, E. 2003, New Features of SAOImage DS9, in Astronomical Society of the Pacific Conference Series, Vol. 295, Astronomical Data Analysis Software and Systems XII, ed. H. E. Payne, R. I. Jedrzejewski, & R. N. Hook, 489
- Juric, M. 2014, mpsky: Multi-purpose sky catalog cross-matching, <https://github.com/mjuric/mpsky>
- Jurić, M., Ciardi, D., Dubois-Felsmann, G., & Guy, L. 2019, LSST Science Platform Vision Document, Systems Engineering Controlled Document LSE-319, NSF-DOE Vera C. Rubin Observatory. <https://lse-319.lsst.io/>
- Jurić, M., Kantor, J., Lim, K. T., et al. 2017, The LSST Data Management System, in Astronomical Society of the Pacific Conference Series, Vol. 512, Astronomical Data Analysis Software and Systems XXV, ed. N. P. F. Lorente, K. Shortridge, & R. Wayth, 279, doi: [10.48550/arXiv.1512.07914](https://doi.org/10.48550/arXiv.1512.07914)
- Jurić, M., Axelrod, T., Becker, A., et al. 2023, Data Products Definition Document, Systems Engineering Controlled Document LSE-163, NSF-DOE Vera C. Rubin Observatory. <https://lse-163.lsst.io/>
- Kahn, S. M., Kurita, N., Gilmore, K., et al. 2010, Design and development of the 3.2 gigapixel camera for the Large Synoptic Survey Telescope, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7735, Ground-based and Airborne Instrumentation for Astronomy III, ed. I. S. McLean, S. K. Ramsay, & H. Takami, 77350J, doi: [10.1117/12.857920](https://doi.org/10.1117/12.857920)
- Kannawadi, A. 2022, Consistent galaxy colors with Gaussian-Aperture and PSF photometry, Data Management Technical Note DMTN-190, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-190.lsst.io/>
- Karavakis, E., Guan, W., Yang, Z., et al. 2024, Integrating the PanDA Workload Management System with the Vera C. Rubin Observatory, in European Physical Journal Web of Conferences, Vol. 295, European Physical Journal Web of Conferences (EDP), 04026, doi: [10.1051/epjconf/202429504026](https://doi.org/10.1051/epjconf/202429504026)
- Knight, S. 2005, Building software with SCons, Computing in Science Engineering, 7, 79, doi: [10.1109/MCSE.2005.11](https://doi.org/10.1109/MCSE.2005.11)
- Krekel, H. 2017, pytest: helps you write better programs, <https://docs.pytest.org>
- Kron, R. G. 1980, Photometry of a complete sample of faint galaxies., ApJS, 43, 305, doi: [10.1086/190669](https://doi.org/10.1086/190669)
- Kuijken, K. 2008, GaaP: PSF- and aperture-matched photometry using shapelets, A&A, 482, 1053, doi: [10.1051/0004-6361:20066601](https://doi.org/10.1051/0004-6361:20066601)
- Kuijken, K., Heymans, C., Hildebrandt, H., et al. 2015, Gravitational lensing analysis of the Kilo-Degree Survey, MNRAS, 454, 3500, doi: [10.1093/mnras/stv2140](https://doi.org/10.1093/mnras/stv2140)
- Labrie, K., Simpson, C., Cardenes, R., et al. 2023, DRAGONS-A Quick Overview, Research Notes of the American Astronomical Society, 7, 214, doi: [10.3847/2515-5172/ad0044](https://doi.org/10.3847/2515-5172/ad0044)
- Li, X., Miyatake, H., Luo, W., et al. 2022, The three-year shear catalog of the Subaru Hyper Suprime-Cam SSP Survey, PASJ, 74, 421, doi: [10.1093/pasj/psac006](https://doi.org/10.1093/pasj/psac006)
- Lim, K.-T. 2022, Proposal and Prototype for Prompt Processing, Data Management Technical Note DMTN-219, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-219.lsst.io/>
- Lupton, R., Plazas Malagón, A. A., & Waters, C. 2025, Verifying LSST Calibration Data Products, Data Management Technical Note DMTN-101, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-101.lsst.io/>
- Lust, N. B., Jenness, T., Bosch, J. F., et al. 2023, Data management and execution systems for the Rubin Observatory Science Pipelines, arXiv e-prints, arXiv:2303.03313, doi: [10.48550/arXiv.2303.03313](https://doi.org/10.48550/arXiv.2303.03313)
- Mandelbaum, R., Hirata, C. M., Seljak, U., et al. 2005, Systematic errors in weak lensing: application to SDSS galaxy-galaxy weak lensing, MNRAS, 361, 1287, doi: [10.1111/j.1365-2966.2005.09282.x](https://doi.org/10.1111/j.1365-2966.2005.09282.x)
- Mandelbaum, R., Miyatake, H., Hamana, T., et al. 2018, The first-year shear catalog of the Subaru Hyper Suprime-Cam Subaru Strategic Program Survey, PASJ, 70, S25, doi: [10.1093/pasj/psx130](https://doi.org/10.1093/pasj/psx130)
- McCormick, J., Dubois-Felsmann, G. P., Salnikov, A., Van Klaveren, B., & Jenness, T. 2024, Using Felis to Represent the Semantics and Metadata of Astronomical Data Catalogs, arXiv e-prints, arXiv:2412.09721, doi: [10.48550/arXiv.2412.09721](https://doi.org/10.48550/arXiv.2412.09721)

- Melchior, P., Joseph, R., & Moolekamp, F. 2019, Proximal Adam: Robust Adaptive Update Scheme for Constrained Optimization, arXiv e-prints, arXiv:1910.10094, doi: [10.48550/arXiv.1910.10094](https://doi.org/10.48550/arXiv.1910.10094)
- Melchior, P., Moolekamp, F., Jerdee, M., et al. 2018, SCARLET: Source separation in multi-band images by Constrained Matrix Factorization, *Astronomy and Computing*, 24, 129, doi: [10.1016/j.ascom.2018.07.001](https://doi.org/10.1016/j.ascom.2018.07.001)
- Merritt, S. R., Fedorets, G., Schwamb, M. E., et al. 2025, Sorchia: A Solar System Survey Simulator for the Legacy Survey of Space and Time, arXiv e-prints, arXiv:2506.02804, doi: [10.48550/arXiv.2506.02804](https://doi.org/10.48550/arXiv.2506.02804)
- Miyazaki, S., Komiyama, Y., Kawanomoto, S., et al. 2018, Hyper Suprime-Cam: System design and verification of image quality, *PASJ*, 70, S1, doi: [10.1093/pasj/psx063](https://doi.org/10.1093/pasj/psx063)
- Morrison, C. B. 2018, Pessimistic Pattern Matching for LSST, Data Management Technical Note DMTN-031, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-031.lsst.io/>
- Mueller, F., et al. 2023, Qserv: A Distributed Petascale Database for the LSST Catalogs, 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>
- Mullaney, J. R., Makrygianni, L., Dhillion, V., et al. 2021, Processing GOTO data with the Rubin Observatory LSST Science Pipelines I: Production of coadded frames, *PASA*, 38, e004, doi: [10.1017/pasa.2020.45](https://doi.org/10.1017/pasa.2020.45)
- NSF-DOE Vera C. Rubin Observatory. 2020, LSST Atmospheric Transmission Imager and Slitless Spectrograph (LATISS), NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2571930](https://doi.org/10.71929/rubin/2571930)
- NSF-DOE Vera C. Rubin Observatory. 2025, Legacy Survey of Space and Time Data Preview 1 [Data set], NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/RUBIN/2570308](https://doi.org/10.71929/RUBIN/2570308)
- Okura, Y., Petri, A., May, M., Plazas, A. A., & Tamagawa, T. 2016, Consequences of CCD Imperfections for Cosmology Determined by Weak Lensing Surveys: From Laboratory Measurements to Cosmological Parameter Bias, *The Astrophysical Journal*, 825, 61, doi: [10.3847/0004-637X/825/1/61](https://doi.org/10.3847/0004-637X/825/1/61)
- Okura, Y., Plazas, A. A., May, M., & Tamagawa, T. 2015, Spurious shear induced by the tree rings of the LSST CCDs, *Journal of Instrumentation*, 10, C08010, doi: [10.1088/1748-0221/10/08/C08010](https://doi.org/10.1088/1748-0221/10/08/C08010)
- O'Mullane, W., Economou, F., Lim, K.-T., et al. 2022, Software Architecture and System Design of Rubin Observatory, arXiv e-prints, arXiv:2211.13611, doi: [10.48550/arXiv.2211.13611](https://doi.org/10.48550/arXiv.2211.13611)
- O'Mullane, W., Economou, F., Huang, F., et al. 2024, Rubin Science Platform on Google: the story so far, 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](https://doi.org/10.48550/arXiv.2111.15030)
- O'Mullane, W., Allbery, R., AlSayyad, Y., et al. 2024, Rubin Observatory Data Security Standards Implementation, Data Management Technical Note DMTN-199, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-199.lsst.io/>
- Padmanabhan, N., Lupton, R., & Loomis, C. 2015, EUPS — a Tool to Manage Software Dependencies, <https://github.com/RobertLuptonTheGood/eups>
- Parejko, J., & Astier, P. 2018, jointcal: Simultaneous Astrometry & Photometry for thousands of Exposures with Large CCD Mosaics, Data Management Technical Note DMTN-036, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-036.lsst.io/>
- Park, H., Karpov, S., Nomerotski, A., & Tsybychev, D. 2020, Tree rings in Large Synoptic Survey Telescope production sensors: its dependence on radius, wavelength, and back bias voltage, *Journal of Astronomical Telescopes, Instruments, and Systems*, 6, 011005, doi: [10.1117/1.JATIS.6.1.011005](https://doi.org/10.1117/1.JATIS.6.1.011005)
- Park, H. Y., Nomerotski, A., & Tsybychev, D. 2017, Properties of tree rings in LSST sensors, *Journal of Instrumentation*, 12, C05015, doi: [10.1088/1748-0221/12/05/C05015](https://doi.org/10.1088/1748-0221/12/05/C05015)
- Paszke, A., Gross, S., Massa, F., et al. 2019, PyTorch: An Imperative Style, High-Performance Deep Learning Library, in *Advances in Neural Information Processing Systems*, ed. H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, & R. Garnett, Vol. 32 (Curran Associates, Inc.), doi: [10.48550/arXiv.1912.01703](https://doi.org/10.48550/arXiv.1912.01703)
- Patterson, M., Bellm, E., Swinbank, J., & Nelson, S. 2020, Design of the LSST Alert Distribution System, Data Management Technical Note DMTN-093, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-093.lsst.io/>
- Pier, J. R., Munn, J. A., Hindsley, R. B., et al. 2003, Astrometric Calibration of the Sloan Digital Sky Survey, *AJ*, 125, 1559, doi: [10.1086/346138](https://doi.org/10.1086/346138)
- Plazas Malagón, A. A., Waters, C., Broughton, A., et al. 2025, Instrument signature removal and calibration products for the Rubin Legacy Survey of Space and Time, *Journal of Astronomical Telescopes, Instruments, and Systems*, 11, 011209, doi: [10.1117/1.JATIS.11.1.011209](https://doi.org/10.1117/1.JATIS.11.1.011209)
- Reed, S. L. 2025, An introduction to Analysis Tools, Data Management Technical Note DMTN-314, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-314.lsst.io/>

- Reiss, D. J., & Lupton, R. H. 2016, Implementation of Image Difference Decorrelation, Data Management Technical Note DMTN-021, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-021.lsst.io/>
- Roby, W., Wu, X., Dubois-Felmann, G., et al. 2020, Firefly and Python — New Ways to Visualize Data on the Web, in *Astronomical Society of the Pacific Conference Series*, Vol. 527, *Astronomical Data Analysis Software and Systems XXIX*, ed. R. Pizzo, E. R. Deul, J. D. Mol, J. de Plaa, & H. Verhouwer, 243
- Roodman, A., Rasmussen, A., Bradshaw, A., et al. 2024, LSST camera verification testing and characterization, in *Ground-based and Airborne Instrumentation for Astronomy X*, ed. J. J. Bryant, K. Motohara, & J. R. D. Vernet, Vol. 13096, *International Society for Optics and Photonics (SPIE)*, 130961S, doi: [10.1117/12.3019698](https://doi.org/10.1117/12.3019698)
- Rowe, B. T. P., Jarvis, M., Mandelbaum, R., et al. 2015, GALSIM: The modular galaxy image simulation toolkit, *Astronomy and Computing*, 10, 121, doi: [10.1016/j.ascom.2015.02.002](https://doi.org/10.1016/j.ascom.2015.02.002)
- Salnikov, A., & McCormick, J. 2024, Current status of APDB and PPDB implementation, Data Management Technical Note DMTN-293, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-293.lsst.io/>
- Saunders, C. 2021, Streak Masking in DM Image Processing, Data Management Technical Note DMTN-197, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-197.lsst.io/>
- Saunders, C. 2024, Astrometric Calibration in the LSST Pipeline, Data Management Technical Note DMTN-266, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-266.lsst.io/>
- Schutt, T., Jarvis, M., Roodman, A., et al. 2025, Dark Energy Survey Year 6 Results: Point-Spread Function Modeling, *The Open Journal of Astrophysics*, 8, 26, doi: [10.33232/001c.132299](https://doi.org/10.33232/001c.132299)
- Sérsic, J. L. 1963, Influence of the atmospheric and instrumental dispersion on the brightness distribution in a galaxy, *Boletín de la Asociación Argentina de Astronomía La Plata Argentina*, 6, 41
- Sérsic, J. L. 1968, *Atlas de Galaxias Australes* (Observatorio Astronómico, Universidad Nacional de Córdoba)
- Shupe, D. L., Moshir, M., Li, J., et al. 2005, The SIP Convention for Representing Distortion in FITS Image Headers, in *Astronomical Society of the Pacific Conference Series*, Vol. 347, *Astronomical Data Analysis Software and Systems XIV*, ed. P. Shopbell, M. Britton, & R. Ebert, 491
- SLAC National Accelerator Laboratory, & NSF-DOE Vera C. Rubin Observatory. 2024, LSST Commissioning Camera, SLAC National Accelerator Laboratory (SLAC), Menlo Park, CA (United States), doi: [10.71929/RUBIN/2561361](https://doi.org/10.71929/RUBIN/2561361)
- SLAC National Accelerator Laboratory, & NSF-DOE Vera C. Rubin Observatory. 2025, The LSST Camera (LSSTCam), SLAC National Accelerator Laboratory (SLAC), Menlo Park, CA (United States), doi: [10.71929/rubin/2571927](https://doi.org/10.71929/rubin/2571927)
- Stalder, B., Reil, K., Aguilar, C., et al. 2022, Rubin Observatory Commissioning Camera: summit integration, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 12184, *Ground-based and Airborne Instrumentation for Astronomy IX*, ed. C. J. Evans, J. J. Bryant, & K. Motohara, 121840J, doi: [10.1117/12.2630184](https://doi.org/10.1117/12.2630184)
- Sutherland, W., Emerson, J., Dalton, G., et al. 2015, The Visible and Infrared Survey Telescope for Astronomy (VISTA): Design, technical overview, and performance, *A&A*, 575, A25, doi: [10.1051/0004-6361/201424973](https://doi.org/10.1051/0004-6361/201424973)
- Tabur, V. 2007, Fast Algorithms for Matching CCD Images to a Stellar Catalogue, *PASA*, 24, 189, doi: [10.1071/AS07028](https://doi.org/10.1071/AS07028)
- Taranu, D. S. 2025, The MultiProFit astronomical source modelling code, Data Management Technical Note DMTN-312, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-312.lsst.io/>
- Thomas, S. J., Barr, J., Callahan, S., et al. 2022, Rubin Observatory Simonyi Survey Telescope status overview, 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, 121820W, doi: [10.1117/12.2630226](https://doi.org/10.1117/12.2630226)
- Tyson, J. A., Ivezić, Ž., Bradshaw, A., et al. 2020, Mitigation of LEO Satellite Brightness and Trail Effects on the Rubin Observatory LSST, *AJ*, 160, 226, doi: [10.3847/1538-3881/abba3e](https://doi.org/10.3847/1538-3881/abba3e)
- van Rossum, G. 2013, PEP 8 – Style Guide for Python Code, Python Software Foundation. <https://www.python.org/dev/peps/pep-0008/>
- Vera C. Rubin Observatory. 2025, The Vera C. Rubin Observatory Data Preview 1, Technical Note RTN-095, NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2570536](https://doi.org/10.71929/rubin/2570536)
- Vereš, P., Jedicke, R., Denneau, L., et al. 2012, Improved Asteroid Astrometry and Photometry with Trail Fitting, *PASP*, 124, 1197, doi: [10.1086/668616](https://doi.org/10.1086/668616)

- Wang, D. L., Monkewitz, S. M., Lim, K.-T., & Becla, J. 2011, Qserv: A Distributed Shared-nothing Database for the LSST Catalog, in *State of the Practice Reports, SC '11* (New York, NY, USA: ACM), 12:1–12:11, doi: [10.1145/2063348.2063364](https://doi.org/10.1145/2063348.2063364)
- Wang, S.-Y., Huang, P.-J., Chen, H.-Y., et al. 2020, Prime Focus Spectrograph (PFS): the prime focus instrument, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 11447, *Ground-based and Airborne Instrumentation for Astronomy VIII*, ed. C. J. Evans, J. J. Bryant, & K. Motohara, 114477V, doi: [10.1117/12.2561194](https://doi.org/10.1117/12.2561194)
- Waters, C., & Rykoff, E. 2025, Calibration Generation, Verification, Acceptance, and Certification., *Data Management Technical Note DMTN-222*, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-222.lsst.io/>
- Waters, C. Z., Magnier, E. A., Price, P. A., et al. 2020, Pan-STARRS Pixel Processing: Detrending, Warping, Stacking, *ApJS*, 251, 4, doi: [10.3847/1538-4365/abb82b](https://doi.org/10.3847/1538-4365/abb82b)
- Wright, A. H., Kuijken, K., Hildebrandt, H., et al. 2024, The fifth data release of the Kilo Degree Survey: Multi-epoch optical/NIR imaging covering wide and legacy-calibration fields, *A&A*, 686, A170, doi: [10.1051/0004-6361/202346730](https://doi.org/10.1051/0004-6361/202346730)
- Zackay, B., Ofek, E. O., & Gal-Yam, A. 2016, Proper Image Subtraction—Optimal Transient Detection, Photometry, and Hypothesis Testing, *ApJ*, 830, 27, doi: [10.3847/0004-637X/830/1/27](https://doi.org/10.3847/0004-637X/830/1/27)
- Zhang, T., Almoubayyed, H., Mandelbaum, R., et al. 2023, Impact of point spread function higher moments error on weak gravitational lensing - II. A comprehensive study, *MNRAS*, 520, 2328, doi: [10.1093/mnras/stac3350](https://doi.org/10.1093/mnras/stac3350)

APPENDIX

A. AMPLIFIER OFFSET CORRECTION

The amplifier offset (amp-offset) correction runs as part of the instrument signature removal (ISR) process. This correction is designed to address systematic discontinuities in background sky levels across amplifier boundaries. It is believed that these discontinuities arise from electronic biases between adjacent amplifiers, persisting even after the application of dark and flat corrections.

Drawing on the **PanSTARRS** Pattern Continuity algorithm (C. Z. Waters et al. 2020), our method aims to eliminate these offsets, thereby preventing problems such as background over-/under-subtraction at amplifier boundaries caused by discontinuities across the detector.

The amp-offset algorithm initially computes a robust flux difference measure between two narrow strips on opposite sides of each amplifier-amplifier interface. Regions containing detected sources, or pixel data which have been masked for other reasons, are not considered. These amp-interface differences are stored in an amp-offset matrix; diagonal entries represent the number of neighboring amplifiers, and off-diagonal entries encode information about the associations between amplifiers. A complementary interface matrix encodes directional information for these associations. Using this information, a least-squares minimization is performed to determine the optimal pedestal value to be added or subtracted to each amp which would reduce the amp-offset between that amplifier and all of its neighboring amplifiers. This method is generalized to support 2D amplifier geometries within a detector, as with LSSTCam, incorporating length-based weighting into the matrices to account for amplifiers that are not square.