

Faro: A framework for measuring the scientific performance of petascale Rubin Observatory data products

Leanne P. Guy¹, Keith Bechtol^{1,2}, Jeffrey L. Carlin¹, Erik Dennihy¹, Peter S. Ferguson², K. Simon Krughoff¹, Robert H. Lupton³, Colin T. Slater⁴, Krzysztof Findeisen⁴, Arun Kannawadi³, Lee S. Kelvin³, Nate B. Lust³, Lauren A. MacArthur³, Michael N. Martinez², Sophie L. Reed³, Dan S. Taranu³, and W. Michael Wood-Vasey⁵

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

²Department of Physics, University of Wisconsin-Madison, Madison, WI 53706, USA

³Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

⁴University of Washington, Dept. of Astronomy, Box 351580, Seattle, WA 98195, USA

⁵Department of Physics and Astronomy, University of Pittsburgh, 3941 O'Hara Street, Pittsburgh, PA 15260, USA

June 24, 2022

ABSTRACT

The Vera C. Rubin Observatory will advance many areas of astronomy over the next decade with its unique wide-fast-deep multi-color imaging survey, the Legacy Survey of Space and Time (LSST).¹ The LSST will produce approximately 20TB of raw data per night, which will be automatically processed by the LSST Science Pipelines to generate science-ready data products – processed images, catalogs and alerts. To ensure that these data products enable transformative science with LSST, stringent requirements have been placed on their quality and scientific fidelity, for example on image quality and depth, astrometric and photometric performance, and object recovery completeness. In this paper we introduce **faro**, a framework for automatically and efficiently computing scientific performance metrics on the LSST data products for units of data of varying granularity, ranging from single-detector to full-survey summary statistics. By measuring and monitoring metrics, we are able to evaluate trends in algorithmic performance and conduct regression testing during development, compare the performance of one algorithm against another, and verify that the LSST data products will meet performance requirements by comparing to specifications. We present initial results using **faro** to characterize the performance of the data products produced on simulated and precursor data sets, and discuss plans to use **faro** to verify the performance of the LSST commissioning data products.

Keywords: Rubin Observatory, Legacy Survey of Space and Time, LSST, Data Management, Metric, Verification

1. INTRODUCTION

Rubin Observatory’s Legacy Survey of Space and Time (LSST), scheduled to begin operations in 2024, will be the most ambitious and comprehensive optical astronomy survey ever undertaken.¹ The LSST will image the entire visible sky twice each week for 10 years, resulting in a dataset that will comprise over 30 trillion observations of 40 billion astronomical sources. The science enabled by Rubin’s LSST will be very broad, ranging from studies of small moving bodies in the solar system to the structure and evolution of the universe as a whole.

Achieving the ambitious science goals of LSST requires an unprecedented survey dataset and a comprehensive understanding of the survey performance. The LSST Science Requirements Document (SRD)² defines a set of “normative” science performance metrics that describe the quality of the LSST data products needed to meet these goals, for example on image quality and depth, astrometric and photometric performance, and object recovery completeness. By defining, measuring, and tracking science performance metrics, we are able to rapidly assess performance, identify changes, problems, or regressions, and take timely action to ensure the survey stays on track to achieve its goals.

As Rubin Observatory moves into the final stages of construction and begins to prepare for operations, significant effort is being placed on understanding the performance of the as-built system and the quality of the science-ready data products it will deliver. `faro`* is an open-source python framework distributed as part of the LSST science pipelines^{3,4} that computes scalar performance metrics on the outputs of the LSST science pipelines. It provides a customizable and easily-extensible set of tools to:

- Generate artifacts to verify the normative science performance metrics detailed in the LSST Science Requirements Document,
- Compute additional non-normative metrics for science validation,
- Monitor performance, carry out regression analysis, and characterization of the LSST science pipelines,
- Providing a “first-look” data quality analysis capability to inform observatory operations.

2. THE FARO FRAMEWORK

Two main components form the basis of the `faro` framework: a collection of `Tasks` that measure specific metric values and a set of base classes that handle data i/o for various types of input data units corresponding to the granularity of metric computation, e.g., per-detector, per-visit, per-patch, per-tract. `faro` builds upon much of the existing infrastructure of the LSST Science Pipelines, in particular, the “Generation 3” middleware, including the Data Butler and Task Framework,⁵ and the LSST verification framework.⁶ The LSST Data Butler (hereafter the “Butler”) provides an abstracted data access interface that is used to read and write data without knowing the details of file formats or locations. The Butler organizes data, both raw and processed, into data repositories. A `Task` is a reusable unit of code in the LSST science pipelines infrastructure used to process data. Each `Task` has a specialized configuration object attached to it and provides a `run()` method that implements the algorithm to execute. Each `faro` metric has an associated `Task` class that implements the mathematical algorithm to compute the scalar metric value on the input data. The LSST Verification Framework, `lsst.verify`, is a framework for making verification measurements in the LSST Science Pipelines. `faro` takes as input catalog data products generated by the LSST Science Pipelines[†] stored in a Butler repository, and computes scalar performance metrics on them. The resulting metric values are persisted in the same Butler repository alongside the input data products, each with an associated searchable data identifier[‡] as `lsst.verify.Measurement` objects. `faro` has adopted a modular and extensible design that can be configured to run any subset of metrics on any subset of data products, e.g., a subset of the real-time metrics computed on a single exposure.

2.1 Analysis Contexts

An analysis context defines the input data unit corresponding to the granularity of metric computation, e.g., per-detector, per-visit, per-patch, or per-tract. `faro` supports metric calculation for multiple analysis contexts; the same mathematical function can be called in a different analysis context to produce a different metric value. For example, the photometric and astrometric repeatability metrics and band-to-band astrometric transformation accuracy metrics are computed at the visit level on single-frame data products, whereas residual PSF ellipticity correlation metrics could be computed at either a per-visit level or on a per-tract level on the coadded data products. For some metrics, only certain analysis contexts make sense, and often the choice of analysis context will be dictated by statistics, e.g., a small 1 deg^2 CI dataset may not contain enough data to compute some metrics on a per-tract level. `faro` allows users to easily define new analysis contexts for their particular science needs.

The following analysis, corresponding to analysis contexts are currently supported in `faro`:

- per-detector source catalogs, i.e., single-visit detections,

*<https://github.com/lsst/faro>

[†]The LSST science pipelines produce catalog data products both in FITS and parquet file format. `faro` works exclusively with parquet formatted data.

[‡]A dict-like identifier for data usable across multiple Butler collections and `DatasetTypes`.

- per-visit source catalogs, i.e., single-visit detections,
- per-patch object catalogs, i.e., coadd detections both per-band and multi-band,
- per-tract object catalogs, i.e., coadd detections, both per-band and multi-band input,
- per-patch matched source catalogs, i.e., set of single-visit detections of the same objects, both single band and multi-band input,
- per-tract matched source catalogs, i.e., set of single-visit detections of the same objects.

2.2 Base Classes

`faro` provides a set of base classes corresponding to the various analysis contexts described above that use the `PipelineTask` framework to build a quantum graph[§] and interact with Butler for data i/o. The base classes abstract away the data i/o, leaving the scientist free to focus on the algorithmic details relevant to their particular analysis. The primary base classes in the `lsst.faro` package are `CatalogMeasurementBaseConnections`, to define the desired i/o, `CatalogMeasurementBaseConfig`, to provide the configuration, and `CatalogMeasurementBaseTask`, to run the algorithm and store the output data in the Butler. Each of these base classes inherits from `MetricConnections`, `MetricConfig`, and `MetricTask`, in the `lsst.verify` package respectively, and adds additional functionality for computing science performance metrics based on Source and Object catalog inputs. Additional subclasses can be easily added following this scheme for new analysis contexts.

2.3 Stages of Metric Computation

Metric computation with `faro` proceeds through three stages:

1. **Preparation:** any intermediate data products that are needed as input to the subsequent measurement step are assembled and persisted in the Butler.
2. **Measurement:** the metric `Task` is run to compute a measurement for each unit of data (i.e., a quantum of processing with a particular `dataId` for the output measurement). Measurements are stored as `lsst.verify.Measurement` objects.
3. **Summary:** a single scalar summary statistic, e.g., a mean or median, is generated from the collection of input measurements computed in the measurement stage. Summary statistics are stored as `lsst.verify.Measurement` objects and persisted in the Butler.

Not all metrics will necessarily have a preparation stage as there may be no additional intermediate data products needed to compute the metric than those generated as part of the pipelines processing. However, all must have a measurement and summary stage.

For example, the photometric repeatability metric, PA1, measures the RMS photometric repeatability of bright non-saturated unresolved point sources in a single filter. During the preparation stage, a matched source catalog is created for each tract and band that matches source detections in individual visits that correspond to the same physical astronomical object[¶]. During the measurement stage, for each tract and band, the matched catalog computed and stored in the preparation stage is loaded into memory and used to compute the RMS scatter of fluxes for a single astrophysical object. In the final summary stage, the measurements for the ensemble of individual tracts computed in the measurement stage are loaded to compute a median summary statistic per band. This final summary statistic characterizes the overall performance for the dataset per band and is stored as an `lsst.verify` object in the Butler.

[§]A data structure that represents a concrete workflow generated from a Pipeline.

[¶]LSST defines an Object as an astrophysical physical object, e.g., a star, galaxy or asteroid, and a Source as a single detection of an astrophysical object in an image. The association of non-moving Sources leads to Objects; the association of moving Sources leads to Solar System Objects.

2.4 Processing Pipelines

Pipelines, specifications for collections of Tasks in the form of YAML files, are used to define the `farø` metrics to be computed together with the detailed execution parameters for metric calculations. The `farø` package provides a set of pre-defined pipelines, defined per analysis context, for many common configurations of computing metrics. `farø` pipelines can be run as an afterburner to the LSST Science Pipelines or can be interspersed with science pipelines processing to compute, as long as the necessary input data products have been computed and are stored in the Butler. For example, all single-visit metrics can already be computed once the single-frame processing steps have completed. This enables us to promptly assess performance and take any corrective action early. Figure 1 shows an example of a `farø` pipeline describing the measurement stage of the computation of three performance metrics using a per-tract matched catalog analysis context. Pipelines can be built hierarchically, with a high-level pipeline calling lower-level pipelines.

```
description: Compute metrics from matched catalogs
tasks:
  PA1:
    class: lsst.farø.measurement.TractMatchedMeasurementTask
    config:
      connections.package: validate_drp
      connections.metric: PA1
      python: |
        from lsst.farø.measurement import PA1Task
        config.measure.retarget(PA1Task)
  PF1_design:
    class: lsst.farø.measurement.TractMatchedMeasurementTask
    config:
      connections.package: validate_drp
      connections.metric: PF1_design_gri
      python: |
        from lsst.farø.measurement import PF1Task
        config.measure.retarget(PF1Task)
        config.measure.threshPA2 = 15.0
```

Figure 1: A `farø` pipeline describing the measurement stage of the computation of three performance metrics per-tract from matched catalogs.

3. ADDING A METRIC

We anticipate that many developers and scientists will want to define their own metrics specifically to address concerns that arise as LSST progresses through the construction phase and into operations. The first step in writing a new metric is to define the analysis context (§2.1) and to review the associated `Connections`, `Config`, and `Task` base classes (§2.2) to understand the in-memory python objects that will be passed to the run method of the metric measurement task, and the configuration options. If the required analysis context and associated base class do not exist, they must first be created. Once the analysis context and base class are defined, the metric can be implemented by creating a measurement `Task`, an instance of `lsst.pipe.base.Task` that will operate on in-memory python objects and compute the metric. Finally, the new metric is added to the appropriate pipeline YAML file so that it can be run. Figures 2 show an example of a non-normative metric `Task` written to compute the number of rows in an input `Source` or `Object` catalog, `NumSourcesTask`

```

class NumSourcesTask(Task):

    ConfigClass = NumSourcesConfig
    _DefaultName = "numSourcesTask"

    def run(self, metricName, catalog, **kwargs):
        self.log.info("Measuring %s", metricName)
        if self.config.doPrimary:
            nSources = np.sum(catalog[self.config._getColumnName(
                "detect_isPrimary")])
        else:
            nSources = len(catalog)
        self.log.info("Number of sources (nSources) = %i" % nSources)
        meas = Measurement("nsrcMeas", nSources * u.count)
        return Struct(measurement=meas)

```

Figure 2: Implementation of a simple task to count the number of Sources or Objects in an input catalog.

4. TRACKING AND VISUALIZATION

We visualize and track the time evolution of `faro` metrics using the LSST metrics dashboard, SQuaSH.⁷ SQuaSH makes use of the InfluxDB^{||} time series database as the backend store for time-stamped metrics and the Chronograf^{**} user interface. InfluxDB is designed to store large volumes of time series data and quickly perform real-time analysis on that data. Chronograf provides dashboards to easily visualize and track metrics stored in InfluxDB. Metrics computed by `faro` on the outputs of the LSST Science Pipelines are stored in the Butler and then pushed to the SQuaSH using the InfluxDB REST API. From the Chronograf UI, the user can query the metric values, aggregate results, and create alerts to measured metrics that change or go out of specification. By tracking metric measurements, we are able to understand trends in the algorithmic performance of the LSST Science Pipelines and ultimately verify that we will meet our requirements. We have configured several dashboards for tracking `faro` metrics computed on various precursor datasets at different cadences, e.g., nightly, monthly. Figure 3 shows the dashboard for `faro` metrics computed on the nightly CI dataset (§5.1.1).

5. FARO APPLICATIONS

5.1 Monitoring Science Pipelines Performance

`faro` is used to monitor the scientific performance of the LSST science pipelines on precursor datasets of different sizes and regularly processed on a variety of cadences.⁸ By tracking metrics computed on the outputs of evolving versions of the science pipelines using a fixed dataset, we can rapidly identify the effect of code or configuration changes on the data products as development progresses during the LSST construction period.

5.1.1 Continuous Integration

The LSST Science pipelines are run nightly and weekly on a small subset of the Hyper Suprime-Cam (HSC) Release Candidate 2 (HSC RC2) dataset^{††} in the LSST Jenkins-based continuous integration (CI) system. This dataset consists of the central 6 detectors for 8 randomly chosen visits in the 5 broad band filters in the HSC COSMOS field and was produced specifically for measuring metrics in a CI environment. `faro` pipelines are run on the outputs in a per-visit or per-detector analysis context. Given the small size of the dataset, there are not enough statistics to compute metrics on a per-tract level. The goal is to track daily the performance of the pipelines and the quality of the data products and to rapidly identify the effect of code or configuration changes on the data products as development progresses during the LSST construction period. Figure 4 shows

^{||}<https://www.influxdata.com/>

^{**}<https://www.influxdata.com/time-series-platform/chronograf/>

^{††}https://github.com/lsst-dm/rc2_subset

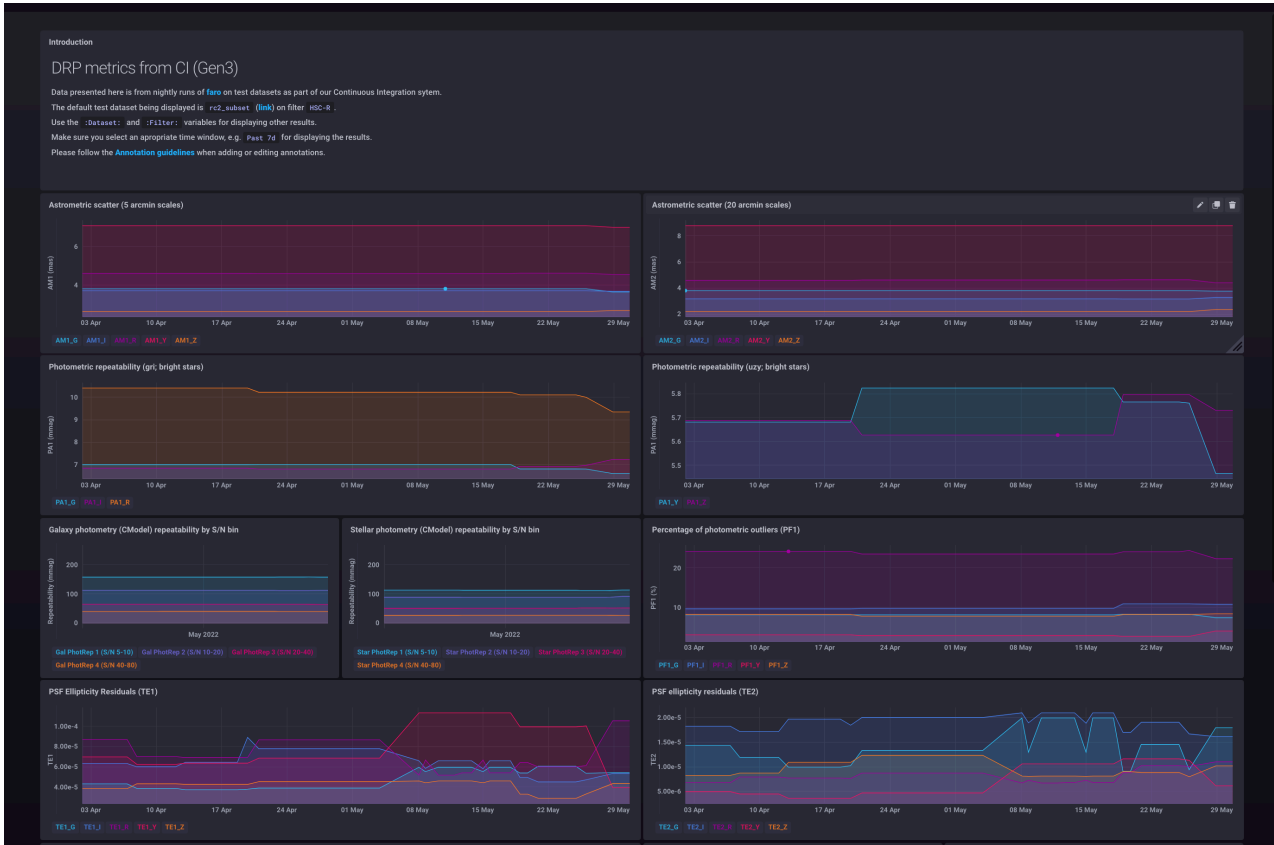
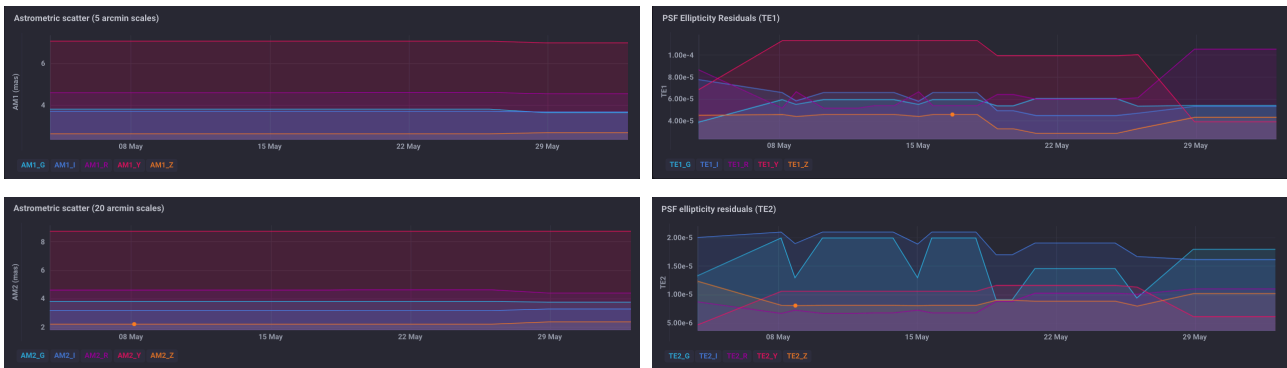


Figure 3: InfluxDB/Chronograph dashboard showing the time evolution of a sample of metrics computed as part of nightly CI on a small subset of the HSC RC2 dataset.

an example of two tracked metrics and how they were used to identify the effect of a software change. Metrics are reviewed regularly allowing us to identify changes and understand the impact of changes in the science pipelines.



(a) Astrometric precision metrics

(b) Residual PSF ellipticity correlation metrics

Figure 4: Examples of the time evolution of four `faro` metrics computed on a nightly cadence in CI over a period of one month on the outputs of the nightly data release processing of a small subset of the HSC PDR2 dataset. Left shows the AM1 metric that characterizes the astrometric scatter on two different spatial scales, 5 and 20 arcmins, and right shows the residual PSF ellipticity correlations, TE1 and TE2.

5.1.2 Software Release Characterization

Major releases of the LSST science pipelines are made approximately once every six months. Every major release is accompanied by a Characterization Metric Report, which describes the scientific performance of the release by computing `far0` metrics on the full HSC RC2 dataset. HSC RC2 consists of 3 tracts of data selected to test various pathological cases, e.g., difficult astrometric solutions, extremely good seeing yielding an under-sampled PSFs, difficult fields for deblending, and large galaxies, among others. These three tracts each contain 112–149 visits split between the HSC-G, HSC-R, HSC-I, HSC-Z, and HSC-Y (grizy) filters. The Characterization Metric Report compares values of metrics computed on the HSC RC2 dataset with a) those computed using the previous release of the science pipelines and b) the LSST SRD design specifications. This allows us to monitor the effects of major strategic developments in algorithms over longer periods. Figure 5 shows an excerpt from the Characterization Metric Report⁹ for the most recent release of the LSST science pipelines, Release 23.0.1. From this excerpt, we see that there were improvements in the photometry and ellipticity correlation metrics as compared to the previous release, which reflect improvements in the photometric calibration algorithms to better handle fields of view near the spatial edges of survey footprints.

Metric	Unit	SRD	Release 22	Release 23	Metric	Unit	SRD	Release 22	Release 23
		Design Value	Value (RC2)	Value (RC2)			Design Value	Value (RC2)	Value (RC2)
PA1: <i>g</i>	mmag	≤ 5.0	7.6	7.1	TE1: <i>g</i>	—	≤ 2 × 10 ⁻⁵	1.5 × 10 ⁻⁵	1.8 × 10 ⁻⁵
PA1: <i>r</i>	mmag	≤ 5.0	8.5	8.4	TE1: <i>r</i>	—	≤ 2 × 10 ⁻⁵	2.6 × 10 ⁻⁴	3.0 × 10 ⁻⁵
PA1: <i>i</i>	mmag	≤ 5.0	9.2	8.7	TE1: <i>i</i>	—	≤ 2 × 10 ⁻⁵	1.9 × 10 ⁻⁵	1.0 × 10 ⁻⁵
PA1: <i>z</i>	mmag	≤ 7.5	7.0	6.7	TE1: <i>z</i>	—	≤ 2 × 10 ⁻⁵	9.4 × 10 ⁻⁶	7.9 × 10 ⁻⁶
PA1: <i>y</i>	mmag	≤ 7.5	8.0	7.9	TE1: <i>y</i>	—	≤ 2 × 10 ⁻⁵	4.8 × 10 ⁻⁵	1.3 × 10 ⁻⁵

Figure 5: Excerpt from the LSST Science Pipelines major Release 23 Characterization Metric Report on the HSC RC2 dataset. Metric values are compared with those of the previous release of the science pipelines and with the design specifications from the LSST Science Requirements document SRD.

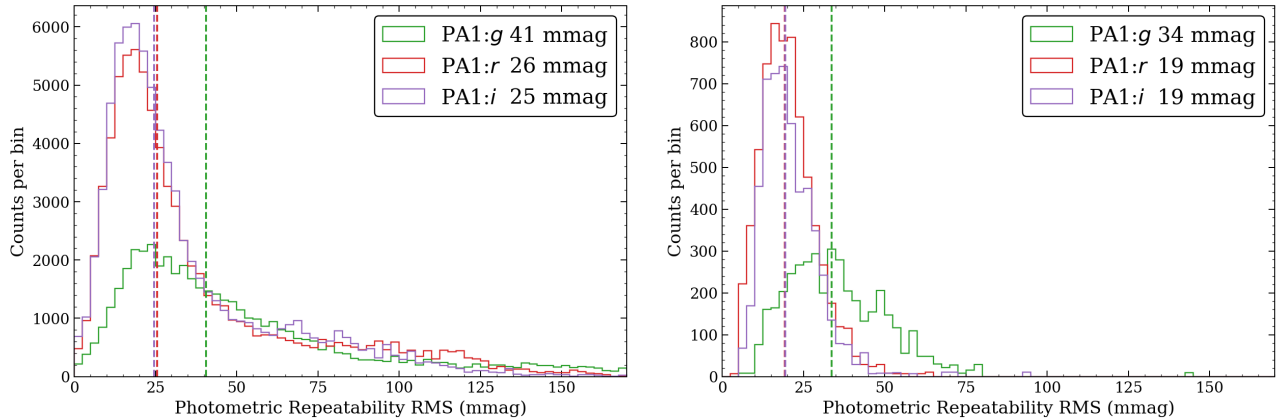
5.2 Rubin Auxiliary Telescope Imaging Surveys

The Rubin Auxiliary Telescope¹⁰ (AuxTel) is a 1.2m f/18 telescope that sits adjacent to Rubin Observatory that will be used to measure atmospheric transmission using broadband spectroscopy of bright stars throughout survey operations. During commissioning, AuxTel is also being used as a control systems and data processing pathfinder via a series of imaging surveys. The imaging surveys are being conducted in three bands (*g*, *r*, and *i*) using a 2x30s exposure sequence per visit. The first AuxTel imaging campaign collected multi-band data over a 1 deg² field of high stellar density in a series of 110 pointings, completing an average of 10 visits per pointing between Feb 2022 and May 2022. The data was processed with the LSST Science Pipelines, including measuring science performance metrics with `far0`. This dataset is sufficient to explore coadds, repeatability, and optimize processing configurations. Single-visit and matched-visit science performance metrics, including the photometric repeatability PA1, were computed with `far0` using data generated from the AuxTel Imaging Survey. Figure 6 shows the PA1 metric values for each bandpass separated into two distributions consider (a) all sources detected and (b) only sources in the center of the detector (1000 < X,Y < 3000 pixels). The physical filters used in the initial imaging campaign suffer from the inward creep of leaked epoxy around their edges, resulting in time-evolving vignetting. This effect is readily seen in the PA1 metric calculation.

5.3 Rapid and first-look analysis

During LSST observing, early analysis of the images coming off the camera will be carried out on the mountain on the timescales of a few minutes. The following are examples of metrics that could be implemented in `far0` to support this rapid first-look analysis:

- Sky brightness, counts, and variance



(a) Photometric repeatability across the full detector (b) Photometric repeatability in the central detector region

Figure 6: The distributions of RMS values of the photometric repeatability of sources detected as part of the initial AuxTel imaging campaign. Vertical dashed lines show the calculated PA1 metric for each bandpass. The left panel shows the value of PA1 calculated for all sources detected and the right panel includes only the sources in the central region of the detector ($1000 < X, Y < 3000$ pixels). The central region of the detector is not affected by the vignetting due the epoxy creep on the filters leading to much tighter distributions and lower PA1 values.

- PSF FWHM and PSF ellipticity
- Comparison to reference catalogs, e.g., Gaia to look for an excess/deficit of detections, astrometric offsets, photometric residuals, or problems with PSF modeling,
- Effective survey depth
- Anomaly detection

5.4 Rubin Data Previews

During the period leading up to the start of operations in 2024, Rubin Observatory will deliver a series of three Data Previews to the community. The goals of these Data Previews are twofold, to serve as an early integration test of the LSST Science Pipelines and the Rubin Science Platform (RSP)^{††, 11} and to enable a limited number of scientists to begin early preparations for science with the LSST. All Rubin Data Previews are hosted at the Interim Data Facility (IDF) on Google cloud.¹² Data Preview 0 (DP0)¹³ is the first of these three data previews. The data set adopted for DP0 is the DESC 300 deg² catalog of simulated LSST-like images and catalogs generated by the Dark Energy Science Collaboration (DESC) for their Data Challenge 2 (DC2).¹⁴ As part of the second phase of Data Preview 0, DP0.2, the DESC DC2 simulated images were reprocessed with Release 23 of the LSST Science Pipelines and released to the community on 7 July 2022. `faro` pipelines were run as part of the DP0.2 reprocessing, interspersed with the DRP pipelines to compute performance metrics. Figure 7 and 8 shows distributions of some of the performance metrics computed as part of DP0.2.

6. CURRENT STATUS AND FUTURE DEVELOPMENT

The `faro` framework is released as part of the LSST Science Pipelines distribution and is running in nightly CI builds on small representative CI datasets and on a few tracts of HSC and DESC DC2 data to track the performance of the LSST Science Pipelines. `faro` is used to produce a detailed Characterization Metric Report

^{††}A set of integrated web applications and services deployed at Rubin Data Access Centers (DACs) through which the scientific community will access, visualize, subset and perform next-to-the-data analysis of LSST Data products.

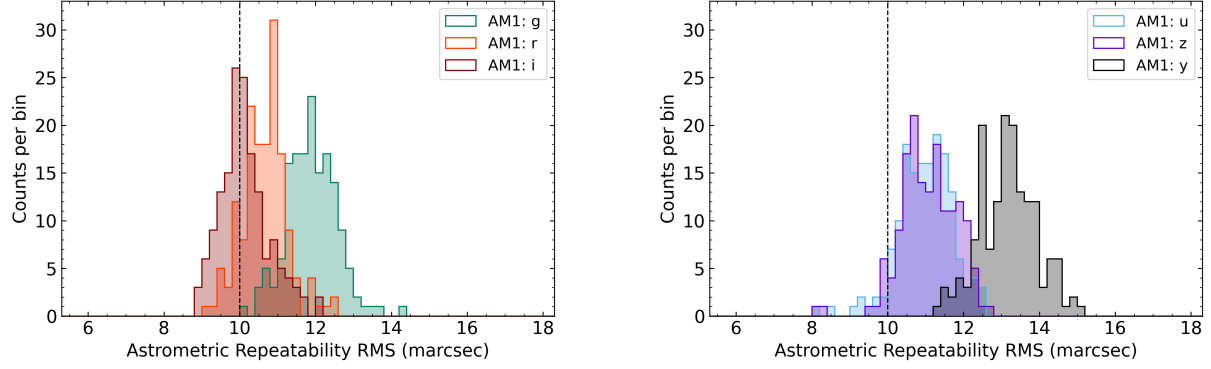


Figure 7: The distributions of the RMS values of the astrometric repeatability on 5' timescales in the a) g , r , and i) and b) u , z , and y bands of sources detected as part of the DP0.2 processing campaign. The vertical dashed lines show the LSST SRD design specifications.




Figure 8: The distributions of the RMS values of the photometric repeatability of bright non-saturated unresolved point sources in the a) g , r , and i) and b) u , z , and l) bands of sources detected as part of the DP0.2 processing campaign. The vertical dashed lines show the LSST SRD design specifications.

for major 6-monthly releases of the pipelines. `faro` is starting to be used to characterize the performance of preliminary data from the Rubin Auxiliary Telescope imaging campaigns and was successfully run at scale on the outputs of the Rubin Operations Data Preview 0.2 processing of the DESC DC2 300 deg² dataset.

Many LSST SRD metrics have been implemented in `faro` spanning a range of analysis contexts; a near-term focus will be on completing the implementation of the normative science performance metrics. The next goal will focus on the needs of LSST commissioning and supporting the implementation of ad-hoc metrics to understand the performance of the as-built system, in particular, to provide near real-time feedback and monitoring of the quality of data taken on the mountain. One indicator of the success of `faro` is the steadily increasing number of active developers defining and implementing metrics to characterize and diagnose performance and issues.

7. CONCLUSIONS

In this paper, we have presented `faro`, the LSST framework for automatically computing scalar performance metrics on the outputs of the LSST science pipelines on a range of spatial and temporal scales. In addition to computing the normative science performance metrics on the LSST data products, `faro` provides a flexible and easy-to-use framework in which anyone can define and write a metric that will be computed on the LSST data products and run as part of the LSST science pipelines. The system is successfully running in nightly and weekly CI on precursor datasets to track the performance of the LSST science pipelines and provide regression analysis as development continues. `faro` has also been successfully demonstrated at scale with the reprocessing of the DESC DC2 dataset¹⁴ at the Rubin Interim Data Facility (IDF) running on Google cloud¹² in the context of Rubin Data Preview 0.2 (DP0.2).¹³ As we enter LSST commissioning, `faro` will be a powerful tool to gain rapid insight into the quality of the data coming off the telescope.

ACKNOWLEDGMENTS

We thank Yusra AlSayyad for her review of this manuscript. This material is based upon work supported in part by the National Science Foundation through Cooperative Agreement AST-1258333 and Cooperative Support Agreement AST-1202910 managed by the Association of Universities for Research in Astronomy (AURA), and the Department of Energy under Contract No. DE-AC02-76SF00515 with the SLAC National Accelerator Laboratory managed by Stanford University. Additional Rubin Observatory funding comes from private donations, grants to universities, and in-kind support from LSSTC Institutional Members.

REFERENCES

- [1] Ivezić, Ž., Kahn, S. M., Tyson, J. A., Abel, B., Acosta, E., Allsman, R., Alonso, D., AlSayyad, Y., Anderson, S. F., Andrew, J., Angel, J. R. P., Angeli, G. Z., Ansari, R., Antilogus, P., Araujo, C., Armstrong, R., Arndt, K. T., Astier, P., Aubourg, É., Auza, N., Axelrod, T. S., Bard, D. J., Barr, J. D., Barrau, A., Bartlett, J. G., Bauer, A. E., Bauman, B. J., Baumont, S., Bechtol, E., Bechtol, K., Becker, A. C., Becla, J., Beldica, C., Bellavia, S., Bianco, F. B., Biswas, R., Blanc, G., Blazek, J., Blandford, R. D., Bloom, J. S., Bogart, J., Bond, T. W., Booth, M. T., Borgland, A. W., Borne, K., Bosch, J. F., Boutigny, D., Brackett, C. A., Bradshaw, A., Brandt, W. N., Brown, M. E., Bullock, J. S., Burchat, P., Burke, D. L., Cagnoli, G., Calabrese, D., Callahan, S., Callen, A. L., Carlin, J. L., Carlson, E. L., Chandrasekharan, S., Charles-Emerson, G., Chesley, S., Cheu, E. C., Chiang, H.-F., Chiang, J., Chirino, C., Chow, D., Ciardi, D. R., Claver, C. F., Cohen-Tanugi, J., Cockrum, J. J., Coles, R., Connolly, A. J., Cook, K. H., Cooray, A., Covey, K. R., Cribbs, C., Cui, W., Cutri, R., Daly, P. N., Daniel, S. F., Daruich, F., Daubard, G., Daues, G., Dawson, W., Delgado, F., Dellapenna, A., de Peyster, R., de Val-Borro, M., Digel, S. W., Doherty, P., Dubois, R., Dubois-Felsmann, G. P., Durech, J., Economou, F., Eifler, T., Eracleous, M., Emmons, B. L., Fausti Neto, A., Ferguson, H., Figueroa, E., Fisher-Levine, M., Focke, W., Foss, M. D., Frank, J., Freeman, M. D., Gangler, E., Gawiser, E., Geary, J. C., Gee, P., Geha, M., Gessner, C. J. B., Gibson, R. R., Gilmore, D. K., Glanzman, T., Glick, W., Goldina, T., Goldstein, D. A., Goodenow, I., Graham, M. L., Gressler, W. J., Gris, P., Guy, L. P., Guyonnet, A., Haller, G., Harris, R., Hascall, P. A., Haupt, J., Hernandez, F., Herrmann, S., Hileman, E., Hoblitt, J., Hodgson, J. A., Hogan, C., Howard, J. D., Huang, D., Huffer, M. E., Ingraham, P., Innes, W. R., Jacoby, S. H., Jain, B., Jammes, F., Jee, M. J., Jenness, T.,

- Jernigan, G., Jevremović, D., Johns, K., Johnson, A. S., Johnson, M. W. G., Jones, R. L., Juramy-Gilles, C., Jurić, M., Kalirai, J. S., Kallivayalil, N. J., Kalmbach, B., Kantor, J. P., Karst, P., Kasliwal, M. M., Kelly, H., Kessler, R., Kinnison, V., Kirkby, D., Knox, L., Kotov, I. V., Krabbendam, V. L., Krughoff, K. S., Kubánek, P., Kuczewski, J., Kulkarni, S., Ku, J., Kurita, N. R., Lage, C. S., Lambert, R., Lange, T., Langton, J. B., Le Guillou, L., Levine, D., Liang, M., Lim, K.-T., Lintott, C. J., Long, K. E., Lopez, M., Lotz, P. J., Lupton, R. H., Lust, N. B., MacArthur, L. A., Mahabal, A., Mand elbaum, R., Markiewicz, T. W., Marsh, D. S., Marshall, P. J., Marshall, S., May, M., McKercher, R., McQueen, M., Meyers, J., Migliore, M., Miller, M., Mills, D. J., Miraval, C., Moeyens, J., Moolekamp, F. E., Monet, D. G., Moniez, M., Monkewitz, S., Montgomery, C., Morrison, C. B., Mueller, F., Muller, G. P., Muñoz Arancibia, F., Neill, D. R., Newbry, S. P., Nief, J.-Y., Nomerotski, A., Nordby, M., O’Connor, P., Oliver, J., Olivier, S. S., Olsen, K., O’Mullane, W., Ortiz, S., Osier, S., Owen, R. E., Pain, R., Palecek, P. E., Parejko, J. K., Parsons, J. B., Pease, N. M., Peterson, J. M., Peterson, J. R., Petravick, D. L., Libby Petrick, M. E., Petry, C. E., Pierfederici, F., Pietrowicz, S., Pike, R., Pinto, P. A., Plante, R., Plate, S., Plutchak, J. P., Price, P. A., Prouza, M., Radeka, V., Rajagopal, J., Rasmussen, A. P., Regnault, N., Reil, K. A., Reiss, D. J., Reuter, M. A., Ridgway, S. T., Riot, V. J., Ritz, S., Robinson, S., Roby, W., Roodman, A., Rosing, W., Roucelle, C., Rumore, M. R., Russo, S., Saha, A., Sassolas, B., Schalk, T. L., Schellart, P., Schindler, R. H., Schmidt, S., Schneider, D. P., Schneider, M. D., Schoening, W., Schumacher, G., Schwamb, M. E., Sebag, J., Selvy, B., Sembroski, G. H., Seppala, L. G., Serio, A., Serrano, E., Shaw, R. A., Shipsey, I., Sick, J., Silvestri, N., Slater, C. T., Smith, J. A., Smith, R. C., Sobhani, S., Soldahl, C., Storr-Lombardi, L., Stover, E., Strauss, M. A., Street, R. A., Stubbs, C. W., Sullivan, I. S., Sweeney, D., Swinbank, J. D., Szalay, A., Takacs, P., Tether, S. A., Thaler, J. J., Thayer, J. G., Thomas, S., Thornton, A. J., Thukral, V., Tice, J., Trilling, D. E., Turri, M., Van Berg, R., Vanden Berk, D., Vetter, K., Virieux, F., Vucina, T., Wahl, W., Walkowicz, L., Walsh, B., Walter, C. W., Wang, D. L., Wang, S.-Y., Warner, M., Wiecha, O., Willman, B., Winters, S. E., Wittman, D., Wolff, S. C., Wood-Vasey, W. M., Wu, X., Xin, B., Yoachim, P., and Zhan, H., “LSST: From Science Drivers to Reference Design and Anticipated Data Products,” *ApJ* **873**, 111 (Mar 2019). 1
- [2] Ivezić, Ž. and The LSST Science Collaboration, “LSST Science Requirements Document,” LPM-17 <https://ls.st/LPM-17> (Jan. 2018). 1
- [3] Bosch, J., AlSayyad, Y., Armstrong, R., Bellm, E., Chiang, H.-F., Eggl, S., Findeisen, K., Fisher-Levine, M., Guy, L. P., Guyonnet, A., Ivezić, Ž., Jenness, T., Kovács, G., Krughoff, K. S., Lupton, R. H., Lust, N. B., MacArthur, L. A., Meyers, J., Moolekamp, F., Morrison, C. B., Morton, T. D., O’Mullane, W., Parejko, J. K., Plazas, A. A., Price, P. A., Rawls, M. L., Reed, S. L., Schellart, P., Slater, C. T., Sullivan, I., Swinbank, J. D., Taranu, D., Waters, C. Z., and Wood-Vasey, W. M., “An Overview of the LSST Image Processing Pipelines,” in [*Astronomical Data Analysis Software and Systems XXVII*], Teuben, P. J., Pound, M. W., Thomas, B. A., and Warner, E. M., eds., *Astronomical Society of the Pacific Conference Series* **523**, 521 (2019). 2
- [4] Bosch, J., Armstrong, R., Bickerton, S., Furusawa, H., Ikeda, H., Koike, M., Lupton, R., Mineo, S., Price, P., Takata, T., Tanaka, M., Yasuda, N., AlSayyad, Y., Becker, A. C., Coulton, W., Coupon, J., Garmilla, J., Huang, S., Krughoff, K. S., Lang, D., Leauthaud, A., Lim, K.-T., Lust, N. B., MacArthur, L. A., Mandelbaum, R., Miyatake, H., Miyazaki, S., Murata, R., More, S., Okura, Y., Owen, R., Swinbank, J. D., Strauss, M. A., Yamada, Y., and Yamanoi, H., “The Hyper Suprime-Cam software pipeline,” *PASJ* **70**, S5 (Jan. 2018). 2
- [5] Jenness, T., Bosch, J. F., Salnikov, A., Lust, N. B., Pease, N. M., Gower, M., Kowalik, M., Dubois-Felsmann, G. P., Lim, K.-T., Mueller, F., and Schellart, P., “The Vera C. Rubin Observatory Data Butler and Pipeline Execution System,” in [*Software and Cyberinfrastructure for Astronomy VII*], *Proc. SPIE* **12189**, 12189–40 (2022). 2
- [6] Sick, J. and Fausti, A., “LSST Verification Framework API Demonstration,” Vera C. Rubin Observatory SQuaRE Technical Note, SQR-019 <https://sqr-019.lsst.io/> (Apr. 2018). 2
- [7] Fausti, A., “The SQuaSH metrics dashboard,” Vera C. Rubin Observatory SQuaRE Technical Note, SQR-009 <https://sqr-009.lsst.io/> (Dec. 2020). 5
- [8] Wood-Vasey, M., Bellm, E., Bosch, J., Carlin, J., Guy, L., Ivezić, Z., MacArthur, L., and Slater, C., “Test Datasets for Scientific Performance Monitoring,” Vera C. Rubin Observatory Data Management Technical Note, DMTN-091 <https://dmtn-091.lsst.io/> (May 2020). 5

- [9] Carlin, J., “Characterization Metric Report: Science Pipelines Version 23.0.0,” Vera C. Rubin Observatory Data Management Test Report, DMTR-351 <https://dmtr-351.lsst.io/> (Jan. 2022). 7
- [10] Ingraham, P., Clements, A. W., Ribeiro, T., Reuter, M. A., Fisher-Levine, M., Hoblitt, J., Lupton, R. H., Thomas, S., Stubbs, C. W., Arndt, K. T., Callahan, N., Claver, C. F., Colleoni, F., Corral, L., Doherty, P., Economou, F., Neto, A. F., Jenness, T., Khine, H., Krughoff, K. S., Mondrik, N., Menanteau, F., Mills, D. J., O’Mullane, W., Reil, K. A., Rivera, M., Stalder, B., Sebag, J., Shipsey, I., Tighe, R., Thornton, A. J., Villalobos, A., and Wiecha, O., “Vera C. Rubin Observatory auxiliary telescope commissioning as a control system pathfinder,” in [*Software and Cyberinfrastructure for Astronomy VI*], Guzman, J. C. and Ibsen, J., eds., **11452**, 124 – 139, International Society for Optics and Photonics, SPIE (2020). 7
- [11] Jurić, M., Ciardi, D., Dubois-Felsmann, G., and Guy, L., “LSST Science Platform Vision Document,” Vera C. Rubin Observatory, LSE-319 <https://lse-319.lsst.io/> (July 2019). 8
- [12] O’Mullane, W., Economou, F., Huang, F., Speck, D., Chiang, H.-F., Graham, M. L., Allbery, R., Banek, C., Sick, J., Thornton, A. J., Masciarelli, J., Lim, K.-T., Mueller, F., Padolski, S., Jenness, T., Krughoff, K. S., Gower, M., Guy, L. P., and Dubois-Felsmann, G. P., “Rubin Science Platform on Google: the story so far,” in [*Astronomical Data Analysis Software and Systems XXXI*], *ASP Conf. Ser.* **in press**, arXiv:2111.15030 (Nov. 2021). 8, 10
- [13] O’Mullane, W., “Data Preview 0: Definition and planning.,” Vera C. Rubin Observatory Technical Note, RTN-001 <https://rtn-001.lsst.io/> (Sept. 2021). 8, 10
- [14] LSST Dark Energy Science Collaboration (LSST DESC), Abolfathi, B., Alonso, D., Armstrong, R., Aubourg, É., Awan, H., Babuji, Y. N., Bauer, F. E., Bean, R., Beckett, G., Biswas, R., Bogart, J. R., Boutigny, D., Chard, K., Chiang, J., Claver, C. F., Cohen-Tanugi, J., Combet, C., Connolly, A. J., Daniel, S. F., Digel, S. W., Drlica-Wagner, A., Dubois, R., Gangler, E., Gawiser, E., Glanzman, T., Gris, P., Habib, S., Hearin, A. P., Heitmann, K., Hernandez, F., Hložek, R., Hollowed, J., Ishak, M., Ivezić, Ž., Jarvis, M., Jha, S. W., Kahn, S. M., Kalmbach, J. B., Kelly, H. M., Kovacs, E., Korytov, D., Krughoff, K. S., Lage, C. S., Lanusse, F., Larsen, P., Le Guillou, L., Li, N., Longley, E. P., Lupton, R. H., Mandelbaum, R., Mao, Y.-Y., Marshall, P., Meyers, J. E., Moniez, M., Morrison, C. B., Nomerotski, A., O’Connor, P., Park, H., Park, J. W., Peloton, J., Perrefort, D., Perry, J., Plaszczynski, S., Pope, A., Rasmussen, A., Reil, K., Roodman, A. J., Rykoff, E. S., Sánchez, F. J., Schmidt, S. J., Scolnic, D., Stubbs, C. W., Tyson, J. A., Uram, T. D., Villarreal, A. S., Walter, C. W., Wiesner, M. P., Wood-Vasey, W. M., and Zuntz, J., “The LSST DESC DC2 Simulated Sky Survey,” *ApJS* **253**, 31 (Mar. 2021). 8, 10

ACRONYMS

Acronym	Description
API	Application Programming Interface
AST	NSF Division of Astronomical Sciences
AURA	Association of Universities for Research in Astronomy
CI	Continuous Integration
DC2	Data Challenge 2 (DESC)
DE	dark energy
DESC	Dark Energy Science Collaboration
DP0	Data Preview 0
DRP	Data Release Production
FITS	Flexible Image Transport System
FWHM	Full Width at Half-Maximum
HSC	Hyper Suprime-Cam
IDF	Interim Data Facility
LPM	LSST Project Management (Document Handle)
LSST	Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope)
LSSTC	LSST Corporation
PDR2	Public Data Release 2 (HSC)

PSF	Point Spread Function
REST	REpresentational State Transfer
RMS	Root-Mean-Square
RSP	Rubin Science Platform
RTN	Rubin Technical Note
SLAC	SLAC National Accelerator Laboratory
SPIE	The international society for optics and photonics
SQR	SQuARE document handle
SRD	LSST Science Requirements; LPM-17
UI	User Interface
URL	Universal Resource Locator
YAML	Yet Another Markup Language
arcmin	arcminute minute of arc (unit of angle)