



Vera C. Rubin Observatory
Data Management

A Roadmap to Photometric Redshifts for the LSST Object Catalog

M. L. Graham, J. Bosch, L. P. Guy,
and the DM System Science Team.

DMTN-049

Latest Revision: 2022-02-04



Abstract

This roadmap guides the Rubin Observatory Data Management (DM) team's efforts to engage with the scientific community of data-rights holders in order to implement and validate one or more existing photometric redshift (photo- z) estimators into the Data Release (DR) processing pipeline, and serve the resulting photo- z data products for the DR Object catalogs.

The DR Object catalog photo- z estimates will initially meet a set of *minimum scientific attributes* and serve the *widest variety of science applications*. DM will select one or more *existing, community-vetted algorithms* to generate the DR photo- z . The evaluation criteria will include technical considerations related to the DM System.

The Rubin science community has a considerable wealth of expertise in generating photo- z catalogs and will be the primary users of the DR Object catalog photo- z data products. Therefore, the evaluation, implementation, and validation of the DR Object photo- z estimator(s) will be a joint venture between the Rubin construction team and the Rubin science community.

This roadmap is a living document to guide this joint venture, and it will evolve over time to incorporate input from the science community. In 2021, community input regarding the evaluation criteria for photo- z estimators, and which kind of photo- z estimator(s) should be selected, was gathered via DM's series of "LSST Photo- z Virtual Forums" and photo- z "Letters of Recommendation" (LOR). This input has been incorporated into § 4 and 5.

Further on in this roadmap DM will facilitate science community participation in a "Photo- z Validation Cooperative" based on commissioning data for the shortlisted estimators (§ 6; dates TBD).

This roadmap includes possible in-kind contributions from international partnerships to Object catalog photo- z for DR1 and beyond, as well as potential processes of improving the DR Object catalog photo- z and/or federating user-generated photo- z catalogs during Rubin operations.

Change Record

Version	Date	Description	Owner name
1.	2017-04-01	Initial release of preliminary investigation.	Melissa Graham
2.	2018-10-16	Edited to align with recent DPDD updates, some of which were based on the recommendations of Version 1 of this document.	Melissa Graham
3.1	2020-11-30	Updated as per Jira ticket DM-6367. Included community input, and revised to be a roadmap for future community engagement.	Melissa Graham
3.2	2020-12-23	Added Section 3: Technical Considerations; revised Appendix C4: Early Science and Section 6: Commissioning; and other minor revisions.	Leanne Guy
3.3	2021-01-01	Updates to Section 3: Technical Considerations.	Jim Bosch, Leanne Guy
4.	2021-01-20	Released, merged to master branch.	Melissa Graham, Jim Bosch, Leanne Guy
4.1	2021-05-01	Updated Evaluation Criteria and the call for Letters of Recommendation to incorporate community input (tickets/DM-29784). Merged to master branch.	Melissa Graham
4.2	2022-01-XX	Updated after the Letters of Recommendation process completed (tickets/DM-31945).	Melissa Graham

Contents

1 Introduction	1
2 Roadmap and Timeline	2
2.1 Photo-z Coordination Group	3
2.2 Federating a User-Generated Photo-z Catalog	4
2.3 International Partnerships (In-Kind Contributions)	5
3 Technical Considerations from the DM System	6
4 Evaluation Criteria	9
4.1 Scientific Utility	9
4.2 Output Data Products	10
4.3 Scientific Performance	11
4.4 Technical Considerations	11
5 Letters of Recommendation	11
5.1 Summary of the Letters Submitted (and DM's Shortlist)	12
6 Photo-z Validation Cooperative	13
6.1 Seeking Photo-z Expertise for the Rubin Commissioning Team	14
6.2 Potential Commissioning Data for Photo-z Validation	16
A Appendix: Examples of Implementation and Validation	20
A.1 An Example Implementation Process	20
A.2 Example Validation Tests	20
A.3 Example Publications that Evaluate Photo-z Estimator Performance	22
B Appendix: Data Products Related to LSST Photo-z	22
B.1 Inputs to Photo-z Estimators	23
B.2 Training and Calibration Data	23
B.3 Output Schema, Access Methods, and Documentation	24
B.4 Storage and Compression	25

C Appendix: Example Use-Cases for LSST Photo-z	26
C.1 Internal DMS Use-Cases	26
C.2 Scientific Use-Cases	27
C.2.1 Dark Energy	27
C.2.2 Time Domain	28
C.2.3 Galaxies	28
C.2.4 Active Galactic Nuclei	28
C.2.5 Clustering	29
C.2.6 Stars, Milky Way, and Local Volume	29
C.2.7 Education and Public Outreach	29
C.2.8 Science Use-Cases Summary	29
C.3 The Science Impacts of a Data Release Without Object Photo-z	29
C.4 Considerations for Maximizing Early Science	30
C.5 A Summary of Scientific Recommendations from the LORs	31
D Original Call for Photo-z Letters of Recommendation	32
D.1 LOR Template	34

A Roadmap to Photometric Redshifts for the LSST Object Catalog

1 Introduction

A *photometric redshift* is an estimate of an object's cosmological redshift which is based on its photometry (e.g., apparent magnitudes in multiple filters) instead of on its spectral features (e.g., emission and absorption lines). Redshift is a key component of many science goals that will be pursued with data from the Legacy Survey of Space and Time (LSST). Since it will be impossible to obtain spectra for the billions of galaxies that LSST will observe, photometric redshift estimates will be necessary.

Typically, photometric redshift estimators either fit template spectra to the observed photometry or match photometry to a training set of galaxies with spectroscopic redshifts. The latter is often done with machine learning codes; hybrid photo-*z* estimators also exist. Some photo-*z* estimators are more appropriate for some science goals than others, due to the quality or type of results they produce (e.g., point estimates, full posterior probability density functions, or redshift distributions in tomographic bins). For this reason, several research groups in the science community are already planning to generate multiple kinds of photo-*z*.

It would be scientifically prohibitive if all users of LSST data had to generate their own photo-*z* estimates, as this is a computationally intensive calculation. Furthermore, it is a requirement that the LSST Data Management System (DMS) compute a photometric redshift for all detected `Objects`, to be stored in the `Object` catalog. Research and development of new photometric redshift algorithms for LSST data is beyond the scope of the Rubin Observatory construction project as it is an active area of current and future LSST research. The high-level plan is that one or more existing photo-*z* estimator(s) be implemented as part of the Rubin Science Pipelines and run at scale, and/or a user-generated photo-*z* catalog be ingested and federated with the `Object` catalog.

DMS-REQ-0046

This roadmap guides the Rubin Observatory DM team's engagement of the scientific community on its path towards a decision on which *community-vetted* photo-*z* estimator(s) will be implemented for the DR `Object` catalog. This is a **living document** that will evolve over time to incorporate input from the science community, which has a considerable wealth of expertise in generating photo-*z* catalogs and will be the primary users of the photo-*z* data product.

2 Roadmap and Timeline

All items in this timeline reference **actions taken by DM**.

2021-02-01: Wrote a summary of this proposed timeline. Advertised it broadly.

2021-02/04: Held three “LSST Photo-z Virtual Forums”, one per month, which focused on ingesting input from the science community regarding the proposed evaluation criteria (§ 4). Advertised the forums broadly.

2021-05-01: Wrote a call for “Letters of Recommendation” from the science community. Advertised it broadly.

2021-05/08: Held four “LSST Photo-z Virtual Forums”, one per month, which focused on supporting the science community’s efforts to prepare “Letters of Recommendation” (§ 5). Advertised the forums broadly.

2021-09-30: Deadline for “Letters of Recommendation” from the science community.

2022-01-30: Incorporated additional scientific criteria from the “Letters of Recommendation” into § 4. Added a summary of the “Letters” and DM’s shortlist to § 5.1.

Before commissioning, identify community members of the shortlisted photo-z teams to join the Rubin Commissioning Team and help to guide implementation and perform science validation for photo-z estimators (see Section 3.2 of Bechtol et al. (SITCOMTN-010)).

During commissioning, DM and the photo-z Commissioning Team members will facilitate community participation in a “Photo-z Validation Cooperative”, which will analyze photo-z estimates from DM’s shortlisted estimators based on processed commissioning data released as Data Previews (DP) 1 and 2 (§ 6). The timeline for this effort is contingent on the commissioning schedule. The science community is welcome to test non-shortlisted estimators with the data previews. At the end of commissioning, based on the results of the “Photo-z Validation Cooperative”, DM will select one or more estimators to generate the photo-z data products for DR1.

During operations, the Data Production team will generate photo-z for Object catalogs using the software and data products delivered by the Construction-era DMS, and make available any and all supporting materials such as documentation or spectral templates. The Rubin Operations team may solicit and collect community feedback on the photo-z performance, and return to any stage of this roadmap to update the attributes, algorithms, implementation, and/or validation of Object catalog photo-z for future data releases. Additionally (or instead of a DMS-provided photo-z), the Rubin Operations team might choose to federate a user-generated photo-z catalog, as described in § 2.2 below.

Several terms are used above, which are clarified here.

- **Advertise broadly** means posting in the “Photo-z Coordination Group” at [Community.lsst.org](https://community.lsst.org) (§ 2.1), sending emails to Science Collaborations and other lists of potentially interested individuals (e.g., authors of papers containing the terms LSST and photo-z), and including in Rubin newsletters.
- The **science community**, in the context of this document, refers to any individuals or groups of data rights holders who plan to use Rubin data products or services for science – in particular, the Object catalog photo-z.
- “**LSST Photo-z Virtual Forums**” will be informal drop-in discussion sessions for Rubin staff and science community members, and held virtually at a variety of times so as to enable people from all timezones to attend at least one.
- **Ingesting input from the science community** means that this document will be updated to include contributions from the science community to (e.g., updating the proposed criteria in § 4). Science community input will primarily be ingested from the “LSST Photo-z Virtual Forums” and from written discussions in the “Photo-z Coordination Group” at [Community.lsst.org](https://community.lsst.org) (§ 2.1).

2.1 Photo-z Coordination Group

To help all groups and individuals from across the Rubin community communicate about photo-z for the LSST throughout this roadmap, a “Photo-z Coordination Group” has been es-

established as the “Photometric Redshifts” sub-category in the “Science” category at [Community.lsst.org](https://community.lsst.org); see [lsst.org/c1o4381](https://community.lsst.org/c1o4381) for more information. This is not a Group that requires anyone to “join”, but rather it is a centralized location for questions, answers, and discussion about photo-z.

All are encouraged to use this virtual space to raise questions or issues, report plans or progress, seek feedback on ideas, post summary notes from relevant meetings, or post topics and replies about any other LSST photo-z related activities. Posts are welcome from anyone, including e.g., DM staff members, Rubin Operations team members, International Programs members involved with in-kind contributions, Science Collaboration members, and anyone interested in science with LSST photo-z. Active communication should help all parties to optimize and synchronize their efforts and reduce redundancy as progress is made along this roadmap to LSST photo-z.

There are many advantages of using [Community.lsst.org](https://community.lsst.org) for the “Photo-z Coordination Group”. The platform is suitable for asynchronous conversations, which is important for our global community; it is easy to both search and browse, which allows fast access to information; topic features include threaded replies and cross-linking to enable discussion; markdown allows for nice formatting and embedding objects like images; and individuals may be ‘mentioned’ to facilitate collaboration. Furthermore, it is an open platform in which anyone may make an account and contribute, and offers user experience enhancements such as notifications.

2.2 Federating a User-Generated Photo-z Catalog

The construction-era DM System delivered to the Rubin Observatory operations team will provide photo-z with a basic level of scientific capability (§ 4). If those photo-z are rendered obsolete by community efforts, and a superior user-generated photo-z catalog is validated (whose creators are willing to share), then this data product could be federated (i.e., ingested and served with the `Object` catalog). Decisions about federating user-generated data products will be made by the operations team.

As a long-term approach to providing DR `Object` catalog photo-z, federating a user-generated catalog might require providing the community team(s) with access to a small (e.g., ~ 10%) amount of a data release weeks to months in advance. This would enable the community team(s) to train and calibrate their photo-z estimator, and would minimize any delay between

data release and federation (which is desired for many science goals; Appendix C.3). Facilitating multiple teams sharing their photo-*z* catalogs might involve hosting a “photo-*z* server” within the Rubin Science Platform, similar to that of the Dark Energy Survey’s Science Portal, an infrastructure for organizing input catalogs, training and running photo-*z* estimators, and evaluating their output¹, as described by Gschwend et al. (2018)).

2.3 International Partnerships (In-Kind Contributions)

As the process for International Partnerships with Rubin Observatory via in-kind contributions evolved, photometric redshifts emerged as an area in which (1) multiple teams propose to provide different types of contributions and (2) optimizing the LSST scientific returns on these contributions requires coordination across the Rubin community: the in-kind teams, Rubin staff, the Science Collaborations, and the broad community of data rights holders.

Thus, the stages of this roadmap and the “Photo-*z* Coordination Group” (§ 2.1) were developed with the in-kind contributions in mind, to create a timeline and communication methods that will assist with the development and integration of the in-kind program contributions related to photo-*z*.

At the time of Version 4 of this document, the proposed contributions included, e.g., individuals’ expertise; scientific analyses; software algorithms, tools, or infrastructure; computational resources; datasets or observing time. The proposals specify the recipients for these contributions as either the Rubin Observatory project (e.g., DM, commissioning) or the science community (e.g., a Science Collaboration, all data rights holders).

Here are a couple of examples to illustrate how in-kind teams would be involved in this roadmap. The participation level would vary between in-kind teams depending on their type of contribution, for example:

- In-kind teams proposing to contribute photo-*z* related directable effort or software packages to Rubin Observatory might be helping DM to set up infrastructure to run the “Photo-*z* Validation Cooperative”.
- In-kind teams proposing to contribute data or software to the Science Collaborations

¹A series of YouTube tutorials about the DES Science Portal are available at <https://www.youtube.com/playlist?list=PLGFEWqwqBauBIYa8H6KnZ4d-5ytM59vG2>.

might have prepared “Letters of Recommendation”, be writing photo-z algorithms, and planning to participate in the “Photo-z Validation Cooperative”.

As with all other groups and individuals interested in LSST photo-z, the in-kind teams should use the “Photo-z Coordination Group” in the Community Forum for open discussions and Q&A about their contributions, when possible (with the exception of things like staffing or other sensitive aspects, of course). Since this is a living document, this section can be updated as the role of the in-kind contributors clarifies and evolves after the review process concludes.

3 Technical Considerations from the DM System

The following provides a set of *initial* technical considerations for potential photo-z estimators to be implemented in the DM System. The goal is to provide the community with some initial technical boundaries and expectations. This topic was open for discussion during the LSST Photo-z Virtual Forum series, and continued community input on the feasibility or challenges of these technical considerations is welcome at any time (§ 2.1). The DM team endeavors to be clear about the nature of these technical issues, but the following is subject to change as the DM System evolves and as feedback from the community is considered. Adherence with these technical considerations is one of the *proposed* evaluation criteria in § 4.

Scalability: The strongest point of consideration for any contributed photo-z estimator will be its ability to run at LSST scale. The LSST Object catalog is projected to contain ≈ 37 billion Objects by the end of the survey; photo-z estimators must be able to efficiently compute a photometric redshift for all Objects in the DR Object catalog². Shortlisted photo-z estimators will be benchmarked as part of the selection process on either a defined pre-cursor dataset or data taken during LSST commissioning. As an indication, the computation of photometric redshifts should take on the order of milliseconds per object and not seconds per object.

Inputs and Outputs: Photo-z estimators should require only catalog-level LSST inputs and not require the use of pixels. Inputs may include all measurements in all filter bands for a given object in the DR Object catalog. The DM team will ensure that all measured quantities needed as inputs to photo-z estimators are computed and included in the Object catalog (See Appendix B.1). This will include making corrections for MW dust using an external dust map

²An estimator will be considered efficient if it runs in a reasonable amount of time on the available compute resources.

prior to passing fluxes to the photo-*z* estimator(s).

The outputs currently foreseen and budgeted for, irrespective of algorithm, are defined in the LSST Data Products Definition Document (DPDD) Jurić et al. (LSE-163), and are presented in Appendix B. Exact schema definitions for the output user-facing data products are stored in the GitHub repository `sdm_schemas` (https://github.com/lstt/sdm_schemas). This repository represents the source of truth for the schema definitions for user-facing data products. A detailed description of schema management in DM is presented in O’Mullane & Slater (DMTN-153).

Storage Constraints: Additional storage may be required for auxiliary input data such as spectral templates, calibration data or training sets (See Appendix B.2). Storage for outputs defined in the the DPDD Jurić et al. (LSE-163) has already been budgeted for by DM and is presented in Appendix B.1.

External Data Sets: The assembly of vetted training sets (e.g., compilations of spectroscopic and many-band photometric redshifts) is a separate but related aspect with its own technical considerations. Any such external data sets that would be required for a photo-*z* estimator to run on the LSST data must exist by the time of commissioning (or the beginning of final data release processing during operations) so that the algorithms can be trained, run, and validated. These external data sets should also be in the public domain and formatted to be ingested by the estimator.

Estimator Training and Iterative Development: Photo-*z* estimators often require training in order to accurately relate the detailed properties and tendencies of the input photometric measurements to external data. This is rarely something that can be fully automated. Training of estimators such that they can be run at scale on the full Object catalog with no human intervention requires domain-specific knowledge and is not within scope of the Rubin Operations team. Data Management will work with the LSST science community to identify teams who will support the short-listed estimators through the “Photo-*z* Validation Cooperative” during commissioning and beyond, into operations. These community-lead teams would take the lead on ensuring that their estimators are provided pre-trained to run within the LSST science pipelines. This should include both software and auxiliary data that represents (and probably condenses) the results of any training procedure.

To construct such a trained estimator, it is usually necessary to train it with photometry that

is representative of the target data release, on fields that include at least some of the spectroscopic or external many-band photometry the estimator is trained against. The plan is to perform preliminary processing of small areas of sky prior to beginning full data release processing, in order to validate the software to be used in that processing. The Object catalog from this processing would be made available to the photo-*z* community teams to provide an opportunity for interactive training and tuning before the estimator and its auxiliary data are finalized.

The time window for this work will be short, however, and coverage of important fields may be missing, incomplete, or not representative of the survey as a whole. Additionally, it cannot be guaranteed that the range of depths and observational conditions needed to optimally train an estimator will be represented in the preprocessed dataset. While the needs of the photo-*z* estimator will play a role in what preliminary processing is performed, these must be balanced with time constraints and our ability to use those processed data products for global validation. Having a photo-*z* estimator that can operate robustly in the presence of these limitations is thus extremely important (and more important, for our purposes, than the quality of the estimator under more ideal conditions).

The infrastructure necessary to pre-train all estimators, as well as technical support to the teams will be provided by the project. Access to these preliminary processed data products will be subject to the Rubin Data Rights Policy as for all Data Release Data Products. RDO-013

Computational Processing Constraints: Photo-*z* estimators are expected to have a low-memory-footprint and to operate only on measurements in the LSST Object catalog (See Appendix B.1). DM has budgeted for currently known algorithms and continues to revise estimates based on benchmarking of the codebase as development progresses, updating estimates for future processing based on current benchmarking results. Shortlisted community-vetted photo-*z* estimators will be similarly benchmarked to enable DM to accurately revise budget estimates. The DM sizing model through the end of the construction project is described in O’Mullane et al. (DMTN-135).

Implementation Language: The Rubin Science Pipelines are written in a combination of C++ and Python. C++ is employed for computationally intensive operations such as direct pixel-level processing as well as for low-level primitive classes and data structures (e.g. images, PSF models, geometric regions), whereas high-level algorithms are written in Python. Most photo-*z* estimators are expected to fall into the latter category. Estimators written in either C++ or

Python should be easy to adapt to run in the LSST science pipelines. The effort required not only to integrate and run a photo-*z* estimator as part of the Rubin Science Pipelines, but also to support it over the long term in the operations era will need to be understood. For further details on Data Management’s development practices, consult the DM Developer Guide LSST Data Management.

4 Evaluation Criteria

The following set of proposed evaluation criteria was used first as a guide for the “Letters of Recommendation” (§ 5), and will also be used by DM to select photo-*z* estimator(s) to generate the Object catalog photo-*z*. These criteria were assembled by the authors and include input from the science community who attended the “LSST Photo-*z* Virtual Forums” or submitted science use-case “Letters of Recommendation” in 2021. These criteria will continue to evolve – and to become more detailed – during the later stages of this Roadmap, which will further clarify what the community needs and what existing photo-*z* estimators can deliver. Some supporting material that provides background motivation for these criteria can be found in Appendices A, B and C. These criteria would also apply to any user-generated catalogs to be federated with the Object table (Appendix 2.2).

4.1 Scientific Utility

- **Meet Basic Science Needs** – The selected photo-*z* estimator(s) should, at minimum, meet the basic science needs of the community. A *basic* science need would not include, for example, photo-*z* of the quality required for major cosmological advances. The development of such advanced photo-*z* algorithms is an active research topic within the Dark Energy Science Collaboration (The LSST Dark Energy Science Collaboration et al., 2018), and is a significant effort which the Rubin Observatory staff should not attempt to replicate.
- **Serve Diverse Science Needs** – The selected photo-*z* estimator(s) should serve as wide a variety of scientists as possible, especially those who would or could not generate custom photo-*z* for themselves (Appendix C). At least one photo-*z* estimator that can return reliable results not just for galaxies but also for AGN (and, e.g., quasars) should be selected.
- **Demonstrated Success** – The selected photo-*z* estimator(s) should have broad com-

munity adoption and demonstrated success with other wide-field optical surveys. It is especially beneficial if those past surveys overlap with the LSST Main Survey, to enable comparisons.

- **Accessibility** – The selected photo-*z* estimator(s) should have publicly accessible documentation (e.g., journal article, website, schema) that has been used by the science community. The selected photo-*z* estimator(s) should have publicly accessible code, inputs, and metadata, so that all of this can be made available to the science community for reproducibility. Its output data products should be straightforward to understand and easy to access (Appendix B.3).

4.2 Output Data Products

- **Photo-*z* Estimates** – The selected photo-*z* estimator(s) should produce data products that include full posterior distribution functions, and for which derived point estimates with statistics (e.g., mode, mean, standard deviation, skewness, kurtosis; Appendix B.3) can be derived.
- **Photo-*z* Uncertainties** – The selected photo-*z* estimator(s) should provide data products with accurate uncertainties. Photo-*z* estimators that can separate sources of error, e.g., calibration of the input data vs. systematics from templates or spec-*z* training sets, provide more information to scientists and should be prioritized.
- **Flag Parameters** – The selected photo-*z* estimator(s) should provide flags parameters. Flags are often type boolean or integer, and they provide an indication of, e.g., algorithm failure modes, result reliability, input abnormality. Flags (and uncertainties, above) are especially useful to help the novice user avoid misinterpreting the results or overestimating their significance.
- **Galaxy Properties** – Photo-*z* estimators that can provide galaxy properties other than photo-*z*, such as stellar mass or star formation rate, would provide more information to scientists and should be prioritized. For template-fit photo-*z* results, an identifier for the best-fit template should be included in the outputs, and those templates should be publicly accessible.

4.3 Scientific Performance

The selected photo- z estimators should, at least initially, meet a minimum quality in order to serve the basic science needs of the community. As described below, it is likely that many existing photo- z estimators can easily meet a fiducial minimum quality. Of course, from this fiducial minimum there is room for optimization. During the “Photo- z Validation Cooperative”, metrics to quantitatively evaluate the shortlisted estimators will be developed and used to inform the decision of which estimator to select to provide the Object catalog photo- z .

- **Proposed minimum performance targets** – Based on the science use-cases in Appendix C, the Object catalog photo- z could have a point-estimate accuracy of $\sim 10\%$ and still meet the basic science needs. The photo- z results should result in a standard deviation of $z_{\text{true}} - z_{\text{phot}}$ of $\sigma_z < 0.05(1 + z_{\text{phot}})$, and a catastrophic outlier fraction of $f_{\text{outlier}} < 10\%$, over a redshift range of $0.0 < z_{\text{phot}} < 2.0$ for galaxies with $i < 25$ mag galaxies. Note that this preliminary proposed performance has been shown to be achievable for simulated LSST data with existing photo- z estimators (e.g., Graham et al., 2018; Schmidt et al., 2020).
- **Robustness** – Photo- z estimators that are demonstrated to perform well with imperfect priors and/or incomplete training sets, especially at low- z , should be prioritized.

4.4 Technical Considerations

The selected photo- z estimator(s) should meet the technical considerations described in Section 3.

Furthermore, the selected photo- z estimator(s) should be able to meet these technical constraints and produce photo- z for the Object catalog by the time of data release, as there are negative scientific impacts of delaying the addition of photo- z to the Object table (§ C.3).

5 Letters of Recommendation

The original call for photo- z “Letters of Recommendation” (LOR) and the template provided with the call have been moved to Appendix D.

5.1 Summary of the Letters Submitted (and DM’s Shortlist)

The Rubin Observatory Data Management (DM) team sincerely thanks everyone who engaged with the PZ LOR process. In total, there were 20 submissions: 19 LOR and 1 “non-LOR” describing DESC’s photo-z activities. Of the 19 LOR, 12 advocated for particular algorithms, 6 presented scientific use cases, and 1 was a notification of future software development (deep probabilistic networks).

Of the 6 LOR describing scientific use cases:

- 1 focused on galaxies
- 2 focused on dark energy
- 3 focused on active galactic nuclei

The scientific recommendations made by these 6 LOR are detailed in Appendix C.5 and have been incorporated into § 4.

Of the 12 LOR advocating for particular algorithms:

- 4 were machine-learning (ML) based codes (GPz; DEmP; PZFlow; DNF)
- 3 were template-fitting (TF) based codes (LePhare; Phosphoros, BPZ)
- 2 were hybrid ML+TF codes (Delight, ML-accelerated hierarchical SPS models)
- 3 were for codes which performed “post-processing” to enhance PZ estimates (e.g., combine PZ estimates, recalibrate PDFs, refine outlier flags)

The LORs for five of the PZ estimators demonstrated that the software was established and would (or would likely) be capable of meeting the scientific, performance, and technical aspects as described in the call for LORs: GPz, DEmP, DNF, LePhare, and BPZ. The LORs for three of the PZ estimators which are still in development also demonstrated that the software would (or would likely) be capable of meeting these aspects: PZFlow, Delight, and Phosphoros.

Thus, these eight PZ estimators comprise DM’s formal initial shortlist for potential codes to generate PZ for the LSST Objects catalog: GPz, DEmP, DNF, LePhare, BPZ, PZFlow, Delight,

and Phosphoros. Given that DM's ability to support the implementation and validation of PZ estimators during commissioning will be limited, the five more established estimators (GPz, DEmP, DNF, LePhare, BPZ) would be prioritized. Further prioritization would be left to the discretion and expertise of the individuals handling the implementation (§ 6.1).

The LOR for the in-development estimator based on ML-accelerated hierarchical SPS models was, due to the newness of the method, unclear about whether it would likely meet the technical aspects defined in § 3, but seems very promising. The three LORs for the post-processing codes established their potential positive impacts. Furthermore, the DM team is aware that additional PZ Estimators, such as those analyzed in Schmidt et al. (2020), would likely be appropriate for the task of generating PZ for LSST Objects: e.g., ANNz2, EAZY, FlexZ-Boost, METAPhoR, SkyNet, and TPZ.

6 Photo-z Validation Cooperative

As the scientific community has a considerable wealth of expertise in generating photo-z catalogs, and will be the primary users of the photo-z data product, they are the best suited to scientifically validate the Object catalog photo-z.

Towards this end the DM team plans to facilitate the "Photo-z Validation Cooperative", during which DM will focus on providing technical support for the implementation of the eight photo-z estimators that were shortlisted based on the community's "Letters of Recommendation" (§ 5; but the community may also consider other estimators).

As the DM team is very limited in terms of the support it can provide for implementation, several spots on the Rubin Commissioning Team are being made available to individuals with relevant expertise who commit to implementing the shortlisted PZ estimators, to using the unreleased commissioning data for early tests, and to helping guide the science community's validation efforts with the released commissioning data (§ 6.1).

During the "Photo-z Validation Cooperative", the science community would focus on:

- assembling training sets and validation metrics,
- helping to implement photo-z estimators with the LSST Science Pipelines,

- using the data previews to validate the estimators' performance, and
- documenting, evaluating, and sharing the results.

These activities will inform DM's selection of photo-z estimator(s) for the Object catalog, but more broadly would also enable the science community to prepare for photo-z science with the LSST and to develop user-generated photo-z catalogs.

A preliminary collection of photo-z validation techniques are presented in Appendix A, and a preliminary collection of training and calibration data in Appendix B. Types of commissioning data that might be most useful for scientific validation of photo-z estimators are described in § 6.2. Documentation regarding the technical aspects of the implementation and validation process is considered beyond the scope of this document and will appear elsewhere in the future.

The "Photo-z Coordination Group" will be the communications hub for the "Photo-z Validation Cooperative".

6.1 Seeking Photo-z Expertise for the Rubin Commissioning Team

It is expected that *most* of the science community participating in the "Photo-z Validation Cooperative" will use the commissioning data that is released as part of Data Preview 1 and 2 (DP1 and DP2), and will not need to formally join the Rubin Commissioning Team and have access to the unreleased commissioning data.

However, as mentioned above the DM team is very limited in terms of the support it can provide for implementation, and furthermore there are some aspects of implementing and validating photo-z estimators which would benefit from a connection with the unreleased commissioning data.

Thus, Rubin Observatory is seeking to identify a small group of ~5-10 individuals who have expertise with photo-z estimators to join the Rubin Commissioning Team.

As a member of the Rubin Commissioning team, these individuals would be responsible for implementing the shortlisted photo-z estimators and leading the science community's validation efforts during DP1 and DP2. As mentioned in § 5.1, the five more established estimators

(GPz, DEmp, DNF, LePhare, BPZ) should be prioritized for implementation, but any further prioritization or the order of implementation would be left to the discretion and expertise of DM and these individuals doing the implementation work.

Example components of the implementation process which these individuals would be committing to are provided below (see also Appendix A.1).

- adapting code language for the LSST environment
- adapting input/output formats for the LSST catalogs
- ensuring I/O meets storage constraints
- ensuring processing meets computational constraints
- assembling and storing external data sets
- testing estimator training with the unreleased commissioning data
- early science validation based on unreleased commissioning data
- optimizing PZ estimator code to ensure scalability

With the exception of items that require access to the unreleased commissioning data, many of the above overlap with activities that the broader science community will engage in as part of the “Photo-*z* Validation Cooperative” using the commissioning data released as Data Previews. The difference is that the PZ-related Commissioning Team members would be committed to implementation and to providing leadership for the community in these tasks.

Eligibility – All Commissioning Team members must have Rubin data rights, as described in Blum & the Rubin Operations Team (RDO-013). For individuals who are part of US/Chilean community, the Commissioning Announcement of Opportunity will be used to guide the process of identifying Commissioning Team members (e.g., Section 4 of Bechtol et al. (SITCOMTN-010)). Individuals who are associated with an approved international program should consult their program managers to confirm the proposed work fits into the allocated contributions.

DM will reach out to the shortlisted photo-*z* estimator teams as one way to identify potential individuals who are interested in taking on this role with the Rubin Commissioning team, and

anyone else from the science community who is interested in this opportunity to do photo-*z* implementation during Commissioning is welcome to contact Melissa Graham (by March 15).

Interested parties will be asked to prepare a short statement to establish specifically what they would contribute as Commissioning Team members (as individuals or teams) and how much time they can commit. As with all Commissioning Team members, the minimum level of sustained effort should be about one day per week of sustained effort over a year (0.2 FTE) or at least one contiguous month of focused effort (0.1 FTE), or some reasonable interpolation between these “sustained” and “focused” thresholds.

To enable these individuals to get started on implementation in advance of commissioning, if desired, DM has arranged for 5–10 accounts in the Rubin Science Platform (RSP) at the Interim Data Facility (IDF) during Data Preview 0 (DP0) to be reserved for them, as described in (Team & the Operations Executive Team, RTN-004).

6.2 Potential Commissioning Data for Photo-*z* Validation

The commissioning surveys and their schedule remain subject to change; the information in this section will be updated in the future as the situation clarifies. At this time, the commissioning schedule comprises 3 phases:

Early System Integration & Test: The first phase of commissioning comprises 3 months of technical integration and test (I&T) with the Rubin Observatory Commissioning Camera (Com-Cam). The originally planned Key Performance Metrics (KPMs) and 20-year Depth Test and Schedule tests during this phase have since been descoped. This phase is expected to deliver modest amounts of science-quality imaging (e.g., few hours to few nights) with which tests of image quality, single-visit depth, astrometry, and photometry might be possible, depending on the quality of the data. Data Preview 1 (DP1) will be based on data taken during this phase and processed with the Rubin Science Pipelines. This data might be useful as preliminary test data during the implementation of the photo-*z* estimators, and help to test the formats of the input and output data, create visualizations for the validation codes, etc.

Full System Integration & Test with LSSTCam: The second phase of commissioning comprises 3 months of technical integration and test (I&T) with the LSST Camera (LSSTCam) — the science instrument for Rubin Observatory. The originally planned Key Performance Metrics (KPMs), 20-year Depth Test during this phase have since been descoped. This phase is ex-

pected to deliver modest amounts of science-quality imaging (e.g., few hours to few nights), with which tests of image quality, single-visit depth, astrometry, and photometry might be possible, depending on the quality of the data.

Science Verification & Validation: The third and final phase of commissioning will be focused on science verification and validation of the integrated system and will include both wide-area surveys and a full 10-year depth survey in selected reference fields for which external imaging and spectroscopy is available. This phase will last for 2 months.

Data Preview 2 (DP2) will be based on data taken during the second and third phases of commissioning and processed with the Rubin Science Pipelines. DP2 is more likely to provide a preliminary data set for photo-*z* validation and possibly some of the needed data for LSST photo-*z* described in Appendix B.2. Covering well-studied fields with DP2 data (e.g., COSMOS) would facilitate photo-*z* validation; note that the effort to collect community input on the commissioning fields is underway³. As of early 2022, there had been several suggestions for fields that overlap reference datasets which would be suitable for the generation and validation of photo-*z* estimates. The Commissioning Team is taking this input into consideration and will work with the “Photo-*z* Coordination Group” to develop and announce the observing strategy for DP2 as soon as they can.

As a side note for LSST users who are interested in doing science with commissioning data, it should not be assumed that any release of data products based on commissioning surveys will include photo-*z* estimates. The release of any photo-*z* catalogs prior to DR1 remains at the discretion of the DM team.

References

[SITCOMTN-010], Bechtol, K., Claver, C., Test, S.I., et al., 2021, *Announcement of Opportunity: Community Engagement with Rubin Observatory Commissioning Effort*, SITCOMTN-010, URL <https://sitcomtn-010.lsst.io/>

Beck, R., Lin, C.A., Ishida, E.E.O., et al., 2017, MNRAS, 468, 4323 (arXiv:1701.08748), doi:10.1093/mnras/stx687, ADS Link

³<https://community.lsst.org/t/community-input-to-the-on-sky-observing-strategy-during-commissioning>

- [RDO-013]**, Blum, R., the Rubin Operations Team, 2020, *Vera C. Rubin Observatory Data Policy*, RDO-013, URL <https://ls.st/RDO-013>
- Carrasco Kind, M., Brunner, R., 2013, TPZ: Trees for Photo-Z, Astrophysics Source Code Library (ascl:1304.011), ADS Link
- Carrasco Kind, M., Brunner, R.J., 2014, MNRAS, 441, 3550 (arXiv:1404.6442), doi:10.1093/mnras/stu827, ADS Link
- Crocce, M., Ross, A.J., Sevilla-Noarbe, I., et al., 2019, MNRAS, 482, 2807 (arXiv:1712.06211), doi:10.1093/mnras/sty2522, ADS Link
- Dahlen, T., Mobasher, B., Faber, S.M., et al., 2013, ApJ, 775, 93 (arXiv:1308.5353), doi:10.1088/0004-637X/775/2/93, ADS Link
- Graham, M.L., Connolly, A.J., Ivezić, Ž., et al., 2018, AJ, 155, 1 (arXiv:1706.09507), doi:10.3847/1538-3881/aa99d4, ADS Link
- Gschwend, J., Rossel, A.C., Ogando, R.L.C., et al., 2018, Astronomy and Computing, 25, 58 (arXiv:1708.05643), doi:10.1016/j.ascom.2018.08.008, ADS Link
- [RTN-011]**, Guy, L.P., Bellm, E., Blum, B., et al., 2021, *Rubin Observatory Plans for an Early Science Program*, RTN-011, URL <https://rtn-011.lsst.io/>
- Hildebrandt, H., Arnouts, S., Capak, P., et al., 2010, A&A, 523, A31 (arXiv:1008.0658), doi:10.1051/0004-6361/201014885, ADS Link
- [LPM-17]**, Ivezić, Ž., The LSST Science Collaboration, 2018, *LSST Science Requirements Document*, LPM-17, URL <https://ls.st/LPM-17>
- [LSE-163]**, Jurić, M., Axelrod, T., Becker, A., et al., 2021, *Data Products Definition Document*, LSE-163, URL <https://lse-163.lsst.io/>
- LSST Data Management, LSST DM Developer Guide, URL <https://developer.lsst.io/>
- Malz, A.I., Marshall, P.J., DeRose, J., et al., 2018, AJ, 156, 35 (arXiv:1806.00014), doi:10.3847/1538-3881/aac6b5, ADS Link
- [DMTN-153]**, O'Mullane, W., Slater, C., 2020, *Schema Management in DM*, DMTN-153, URL <https://dmtn-153.lsst.io/>
- [DMTN-135]**, O'Mullane, W., Dubois, R., Butler, M., Lim, K.T., 2021, *DM sizing model and cost plan for construction and operations.*, DMTN-135, URL <https://dmtn-135.lsst.io/>

Polsterer, K.L., D'Isanto, A., Gieseke, F., 2016, arXiv e-prints, arXiv:1608.08016 (arXiv:1608.08016), ADS Link

Sánchez, C., Bernstein, G.M., 2019, MNRAS, 483, 2801 (arXiv:1807.11873), doi:10.1093/mnras/sty3222, ADS Link

Sánchez, C., Carrasco Kind, M., Lin, H., et al., 2014, MNRAS, 445, 1482 (arXiv:1406.4407), doi:10.1093/mnras/stu1836, ADS Link

Schmidt, S.J., Malz, A.I., Soo, J.Y.H., et al., 2020, MNRAS, 499, 1587 (arXiv:2001.03621), doi:10.1093/mnras/staa2799, ADS Link

Tanaka, M., Coupon, J., Hsieh, B.C., et al., 2018, PASJ, 70, S9 (arXiv:1704.05988), doi:10.1093/pasj/psx077, ADS Link

[RTN-004], Team, T.C.E., the Operations Executive Team, 2022, *Guidelines for Community Participation in Data Preview 0*, RTN-004, URL <https://rtn-004.lsst.io/>

The LSST Dark Energy Science Collaboration, Mandelbaum, R., Eifler, T., et al., 2018, arXiv e-prints, arXiv:1809.01669 (arXiv:1809.01669), ADS Link

A Appendix: Examples of Implementation and Validation

A.1 An Example Implementation Process

For reference, the basic steps for implementing a photo-z estimator include:

- prepare the LSST data inputs to the photo-z estimators (Appendix B.1)
- prepare the training and calibration data (Appendix B.2)
- install the needed codes (a technical aspect to be described elsewhere)
- run the estimator (on its own and/or embedded in the data release pipeline)
- validate the photo-z estimator outputs (below, A.2)
- store the results in the Object catalog (Appendices B.3 and B.4)
- ensure output schema and access methods are documented (Appendix B.3)

A.2 Example Validation Tests

For reference, this section describes some photo-z validation tests, some of which overlap with the evaluation criteria described in § 4. Validation tests and quality assessment diagnostics will be necessary to ensure the results meet performance expectations (as defined in § 4). These lists are based in part on a brainstorming session during the LSST Project and Community Workshop's session on Photometric Redshifts on Aug 14 2019⁴.

Journal articles that demonstrate validation processes for photo-z from multi-band wide-area surveys include *"DES science portal: Computing photometric redshifts"* (Gschwend et al., 2018); *"Photometric redshifts for Hyper Suprime-Cam Subaru Strategic Program Data Release 1"* (Tanaka et al., 2018); and *"On the realistic validation of photometric redshifts"* (Beck et al., 2017).

Truth Comparisons – Catalogs of true redshifts can be obtained by withholding some fraction of the training set or by cross-matching to external spectroscopic catalogs. Validation metrics should include at least those used to define the minimum performance of the LSST photo-z for basic science needs (§ 4). A list of potential metrics might include the following, and targets or limits on these metric values might apply to subsets in magnitude, color, or redshift:

E.g., Slide 14 of <https://docs.google.com/presentation/d/7fGEahvDQXfJSL41LVjD1ZHV5zpxLhg0rHwvtucs72Ajk/edit?usp=sharing>

- plots of z_{true} vs. z_{phot} for visual inspection
- standard deviation and bias in $\Delta z = (z_{\text{true}} - z_{\text{phot}})/(1 + z_{\text{true}})$
- fraction of (catastrophic) outliers, e.g., $\Delta z > 3\sigma$ or > 0.06 (> 1)
- quantile-quantile (q-q) plots evaluate the shape of $P(z)$
- probability integrated transform, $PIT(z_{\text{phot}}) = \int_0^{z_{\text{phot}}} P(z) dz$, calculated for all test galaxies, should be flat for well-estimated $P(z)$ (Polsterer et al., 2016)
- the continuous ranking probability score, CRPS, should be a lower value for better estimated $P(z)$, as described in (Polsterer et al., 2016)
- redshift confidence for point estimates, $C(z_{\text{phot}}) = \int_{z_{\text{phot}}-0.03}^{z_{\text{phot}}+0.03} P(z) dz$
- loss and risk parameters that characterize the photo- z estimates:
 $L(\Delta z) = 1 - \left(1 + \left(\frac{\Delta z}{\gamma}\right)^2\right)^{-1}$, where γ is a characteristic threshold (e.g., 0.15), and:
 $R(z_{\text{phot}}) = \int P(z) L(z_{\text{phot}}, z) dz$, as described in Tanaka et al. (2018)
- the conditional density estimation (CDE) loss as described in Section 4.2 of Schmidt et al. (2020). (Note that the CDE does not strictly require the true posterior be known.)

Basic Checks– These are tests that can be done using all LSST Objects and do not require a truth catalog.

- the uncertainty in z_{phot} should correlate with photometric error
- the star/galaxy flag parameter should agree with photo- z (stars have $z_{\text{phot}} = 0$)
- evolution of $N(z)$ to higher- z for samples with fainter magnitudes, redder colors

Scientific Applications – The following might not be included in the formal validation process, but represent analyses that the broader scientific community may want to do to inform their use of the LSST photo- z .

- assessing the performance of point estimates in broker photometric classification algorithms
- evaluating the absolute magnitude distribution of Type Ia supernovae once the distance derived from the photo- z point estimates are applied

- evaluating the distributions of derived physical parameters for galaxies using the $P(z)$
- checking whether high SFR galaxies (and maybe other sensitive populations) have reasonable $P(z)$
- galaxy cluster membership identification
- tomographic bin analysis, as in (Crocce et al., 2019)

A.3 Example Publications that Evaluate Photo- z Estimator Performance

- Hildebrandt et al. (2010) tested 18 different photo- z codes on the same sets of simulated and real data and found no significantly outstanding method.
- Dahlen et al. (2013) tested 11 different photo- z codes on the CANDLES data set (U -band through infrared photometry) and also find that no method stands out as the “best”, and that most of the photo- z codes underestimate their redshift errors.
- Sánchez et al. (2014) used the science verification data (200 square degrees of $grizY$ photometry to a depth of $i_{AB} = 24$ magnitudes) of the Dark Energy Survey (DES) to evaluate several photometric redshift estimators. They found that the Trees for Photo- z code (TPZ; Carrasco Kind & Brunner (2013)) provided the most accurate results with the highest redshift resolution, and that template-fitting methods also performed well – especially with priors – but that in general there was no clear “winner.”
- Tanaka et al. (2018) provides a comparative analysis of several photo- z estimators applied to their data set from the HSC Strategic Program. Their website⁵ provides comparative analysis plots for each of them.
- Schmidt et al. (2020) statistically compare the posterior probability distribution functions produced by 12 photo- z estimators for a mock data set that is representative of the LSST, identifying some biases and shortfalls in both the produced PDFs and the evaluation methods used to analyze them.

B Appendix: Data Products Related to LSST Photo- z

For reference, this section provides a **preliminary and incomplete** collection of inputs, outputs, and associated products related to generating and serving LSST photo- z :

⁵<https://hsc-release.mtk.nao.ac.jp/doc/index.php/photometric-redshifts/>

- the LSST data products needed as input to photo- z estimators (Appendix B.1)
- the required training and calibration data, LSST or external (Appendix B.2)
- the outputs' format, access methods, and documentation (Appendix B.3)
- options for compression to store the results of multiple estimators (Appendix B.4)

B.1 Inputs to Photo- z Estimators

It is important to ensure that all measured quantities needed by photometric redshift estimators are going to be computed and included in the `Object` table. Aside from the fluxes and/or apparent magnitudes and errors for each Rubin Observatory filter, which will be provided in the `Object` catalog, the color properties in the `Object` table might be used for photo- z . *"Colors of the object in 'standard seeing' (for example, the third quartile expected survey seeing in the i band, ~ 0.9 arcsec) will be measured. These colors are guaranteed to be seeing-insensitive, suitable for estimation of photometric redshifts"* [LSE-163]. In the `Object` table the relevant elements are:

- `stdColor` (`float[5]`) = *'standard color', color of the object measured in 'standard seeing', suitable for photo- z*
- `stdColorErr` (`float[5]`) = *uncertainty on `stdColor`*

Additionally, measured quantities such as the galaxy size, shape, radial profile, 'clumpiness', or surface brightness; the DCR correction (or residual); or a parameter that represents the clustering density within some radius (e.g., 2 Mpc) might all be useful (e.g., as priors) for photo- z estimators. The effective transmission function (ϕ ; Eq. 5 in LPM-17), which will be provided for all Sources either in the catalog or as a link, is another useful quantity for photo- z estimators.

B.2 Training and Calibration Data

Spectroscopic (or Many-Band) Redshifts – Deep multi-band LSST photometry for spectroscopic fields like COSMOS, and/or WFD-depth LSST photometry that overlaps multiplexed spectroscopic surveys like DESI/4MOST, which is obtained either during commissioning or early in Operations year 1, will likely be necessary to produce photo- z for DR1. Alternatively, redshifts based on many-band photometry can be used.

Wide Area Imaging – Some photo- z methods have requirements other than spec- z fields: e.g., Sánchez & Bernstein (2019) use clustering information to obtain photo- z and this requires wider, shallower field coverage and not a single deep pointing like a spec- z field would have. This wider area would also serve to reduce cosmic variance in the training set (~100 square degrees would serve to average out the variance).

Preliminary LSST Data – For the community to participate in the training or calibration of a photo- z estimator prior to a data release, it will probably be necessary to make a small but representative amount of pre-processed LSST data available in advance of each data release during Operations. (See (§ 3).

B.3 Output Schema, Access Methods, and Documentation

Output Schema – The format of the photo- z output is one of the attributes that must still be defined (§ 4). The photo- z outputs must provide all the necessary inputs to the validation tests, to be defined in Appendix A.2.

The LSST Data Products Definitions Document (DPDD) [LSE-163] defines the format of the Object catalog's table columns which could store the results of photometric redshift estimates, regardless of how they're generated. The following is from Table 5 of the DPDD:

- photoZ (float[2x95]) = photometric redshift likelihood samples – pairs of redshift and likelihood ($z, \log L$) – computed using a to-be-determined published and widely accepted estimator at the time of LSST Commissioning
- photoZ_pest (float[10]) = point estimates for the photometric redshift provided in photoZ

The exact point estimate quantities stored in the photoZ_pest are to-be-determined, *“but likely candidates are the mode, mean, standard deviation, skewness, kurtosis, and 1%, 5%, 25%, 50%, 75%, and 99% points from cumulative distribution”* [LSE-163]. Flags that represent potential catastrophic outliers, failure modes, a photo- z consistent with $z = 0$, etc., could also be included.

Access Methods – The user experience is one of the proposed selection criteria for the LSST photo- z estimator (§ 4). Some examples of publicly released photo- z catalogs which were

prepared with a user experience that might be desirable for the LSST photo- z include the Dark Energy Survey's Science Portal to serve photometric redshifts Gschwend et al. (2018) and the Hyper SuprimeCam Subaru Strategic Program Tanaka et al. (2018)⁶.

If the LSST photo- z are not made available in either the `Objects` table or in a federated or joinable catalog – for example in the case where a community-generated photo- z catalog is replacing the DMS-generated catalog (Appendix 2.2) – and are instead made available via, e.g., a “photo- z server” (as in Gschwend et al. (2018)), then at least the `Object` catalog ID of the most recent data release should be a queryable parameter.

If the results of multiple estimators are generated, compressed, and stored in the `Objects` table, then decompression should be straightforward for the user (Appendix B.4).

Documentation – Appropriate types of documentation might include published journal articles, GitHub repositories, websites, or other online documentation resources (e.g., <https://readthedocs.org/>). Whatever the format, the documentation contents should include:

- general description of the estimator
- adaptations made to ingest LSST data (compared to past applications)
- an analysis of the training, calibration, and validation processes
- a full list of all inputs and outputs (i.e., a schema browser)

B.4 Storage and Compression

The stored values related to photo- z are subject to the storage space allotted in the `Objects` table as described in § B. Both the posteriors and point estimates from several different photo- z estimators could be compressed and stored in this allotted space. Given the variety of use-cases and the fact that different photo- z estimators produce different results (Schmidt et al., 2020), the option to compute, compress, and store estimates from multiple estimators in the 2×95 float might be scientifically desirable.

Efficient $P(z)$ compression algorithms are in development, such as Carrasco Kind & Brunner (2014) and Malz et al. (2018). Carrasco Kind & Brunner (2014) present an algorithm for sparse

⁶<https://hsc-release.mtk.nao.ac.jp/doc/index.php/photometric-redshifts/>

representation, for which “an entire PDF can be stored by using a 4-byte integer per basis function” and “only ten to twenty points per galaxy are sufficient to reconstruct both the individual PDFs and the ensemble redshift distribution, $N(z)$, to an accuracy of 99.9% when compared to the one built using the original PDFs computed with a resolution of $\delta z = 0.01$, reducing the required storage of two hundred original values by a factor of ten to twenty.” Malz et al. (2018) presents a Python package for compressing one-dimensional posterior distribution functions (PDFs), demonstrates its performance on several types of photo- z PDFs, and provides a set of recommendations for best practices which should be consulted when DM is making decisions on the DR photo- z data products.

However, compression (and decompression by users) will require extra computational resources, which should be estimated and considered, and decompression must be fast and easy for users.

C Appendix: Example Use-Cases for LSST Photo- z

This section contains an **incomplete, non-exhaustive** summary of a variety of internal and scientific use-cases for the LSST Object catalog photo- z . These use-cases inform the evaluation criteria proposed in § 4.

Some of the following information on use-cases was collected from participants of the LSST Project and Community Workshop’s session on Photometric Redshifts on Aug 14 2019⁷.

C.1 Internal DMS Use-Cases

DM’s galaxy photometry outputs are being developed with the goal of feeding photometric redshift estimators, so the computation of photometric redshifts is likely to be a part of the science validation process for LSST photometry. Unlike stars, color-color and color-magnitude diagrams for galaxies do not have sufficient structure to reveal issues with the photometry. While other photometric validation techniques will also be useful (such as evaluating the width of galaxy cluster red-sequences) they may only apply to *some* galaxies, whereas *all* galaxies have a redshift.

⁷Thanks to Sam Schmidt, Chris Morrison, Sugata Kaviraj, Gautham Narayan, Lauren Corlies, Travis Rector, Tina Peters, Alex Malz, Dara Norman, Stephen Smartt, and other participants from the science community.

The internal use-case of scientifically validating the galaxy photometry outputs is likely to require a simple photo- z estimator which fits SED templates, since the goal is to evaluate whether the photometric outputs match the colors of real galaxies. Whether such a simple SED-fit photo- z could also serve the scientific use-cases is undetermined, because the photometric validation process is not yet defined or written. The potential internal use-case of needing photo- z in order to assess catalog completeness for low- and high-redshift objects is also likely to be served by a simple SED-fit photo- z estimate.

As a side note, although the DMS will assign fiducial spectral energy distributions (SEDs) to objects in order to apply sub-band wavelength-dependent photometric calibration and PSF modeling, computing photo- z is not planned to be a part of this process. Furthermore, the SED templates used will likely be simpler (e.g., step-function or slope) than would be needed for deriving photo- z .

C.2 Scientific Use-Cases

A variety of potential scientific applications for the object photo- z are discussed in turn. These scientific use-cases should be used to inform the evaluation criteria proposed in § 4. A summary of the commonalities between science use-cases for photo- z is provided in § C.2.8.

C.2.1 Dark Energy

Extragalactic astrophysics such as weak lensing, baryon acoustic oscillations, and Type Ia supernova cosmology are all main science drivers for the LSST, and all require catalogs of galaxies with photometric redshifts. The photo- z estimators for precision cosmology will be custom-tailored to these particular science goals, and the photo- z results are subject to established science requirements for dark energy cosmology (The LSST Dark Energy Science Collaboration et al., 2018). For example, weak lensing and large scale structure require ensemble measurements of $N(z)$ and thus require a full posterior PDF, whereas point-estimate photo- z for individual objects are required for Type Ia supernova host galaxies and the identification of strong lensing candidates and galaxy cluster members. The Dark Energy Science Collaboration (DESC) is developing specialized photo- z pipelines for these science goals (which *could* serve to generate photo- z for the object catalog, as discussed in Appendix 2.2).

C.2.2 Time Domain

The Transients and Variable Stars Science Collaboration reported that they would use LSST-provided `Object photo- z` to identify and/or characterize extragalactic transient host galaxies. Alert packets provide `Object` IDs for the three nearest stars and three nearest galaxies in the most recent data release. Alert stream brokers intend to query the `Objects` catalog in real time to obtain host `photo- z` because photometric classification for transient light curves is *significantly* aided by redshift estimates. The `Object` catalog's `photo- z` will also be used to identify and prioritize the potential host galaxies of gravitational wave events for imaging searches of the optical counterpart.

C.2.3 Galaxies

The Galaxies Science Collaboration reported that they would use LSST-provided `Object photo- z` , and that their science goals require that `photo- z` be accurate enough ($< 10\%$) to derive intrinsic galaxy properties like mass and star formation rate (SFR). They also indicated that posteriors delivered as $P(z, M)$ and/or with rest-frame apparent magnitudes would be useful to their science goals. This indicates that the results of a template-fitting `photo- z` estimator might be more relevant to Galaxies studies than machine-learning estimates (especially if the SED templates are associated with intrinsic galaxy properties like mass, metallicity, or star formation rate). The `Object photo- z` might also be used to assist with star-galaxy separation, to enable population studies, to estimate environmental (clustering) parameters, and/or to choose instrument configurations for spectroscopic follow-up (i.e., the expected location of emission lines).

C.2.4 Active Galactic Nuclei

It is currently unclear how useful the `Object photo- z` will be for the AGN community because there is no special deblending planned for the DMS to produce galaxy photometry which is free of AGN emission. The AGN contribution to the DR CoAdd image stacks, and thus the `Object` catalog photometry, will be an average flux over the LSST survey images. Photometric redshift codes will either have to be able to recognize and deal with AGN contamination, or the `photo- z` estimates for AGN host galaxies will be impacted. Potential AGN contamination could be identified by identifying `DIAObjects` in the nuclear region, but quantifying and removing that AGN flux from the galaxy photometry and recalculating `photo- z` remain a user-generated

data product.

C.2.5 Clustering

Photometric redshifts would likely be used by individuals studying large scale structure and galaxy clustering – for example, as a way to make an initial selection of cluster members.

C.2.6 Stars, Milky Way, and Local Volume

LSST-provided Object photo-z could be used to reject compact extragalactic objects from stellar samples for population studies and/or spectroscopic follow-up campaigns.

C.2.7 Education and Public Outreach

The question “*how far away is it?*” is common to many EPO initiatives and the Object catalog photo-z will be used when preparing information for the public. EPO might also use photo-z for, e.g., generating 3D graphics that visualize large volumes, or educational programs on the Hubble constant. For EPO purposes, high precision is not as important as outlier reduction for photo-z.

C.2.8 Science Use-Cases Summary

Aside from the specialized use-cases related to dark energy cosmology, which will be served by customized photo-z estimators developed within DESC, most other scientific scenarios use the Object photo-z as point estimates of distance in order to subset the data and identifying targets of interest for follow-up, and/or infer intrinsic galaxy properties.

C.3 The Science Impacts of a Data Release Without Object Photo-z

The option to federate a community-generated photo-z catalog leads to this question: *If there will likely be a superior community photo-z anyway, should Data Management avoid installing a photo-z estimator in the DMS, and instead simply wait for the community to generate a catalog?* There are several significant risks and drawbacks to this option.

- An Object table without photo-z at the time of data release is a problem for brokers,

unless alerts are instead (or additionally) associated to an older DR's Object catalog that has photo-*z*. Brokers require host-galaxy photo-*z* to optimally classify and prioritize transients for follow-up, and plan to obtain this information via each alert's associations with nearby Objects from *the most recent DR*.

- The Object photo-*z* might be tailored to the specific science case of the community team and might not serve the broader science use-cases.
- There is no incentive or reward for a community team that generates a photo-*z* catalog, except perhaps citations to their photo-*z* catalog. Consequently, if no community team generates and donates a catalog there would be no Object photo-*z*, which is a risk for the science use-cases described in Appendix C.2

C.4 Considerations for Maximizing Early Science

Plans for early science with LSST are still in development and are strongly dependent on the outcomes of commissioning. Given the uncertainties on the actual construction schedule and commissioning period, several different plans for early science are being considered (RTN-011). The actual state of science verification (SV) and system completeness at handover to Operations will determine the most appropriate course of action to take regarding early science.

In the event that SV is completely successful and the project moves quickly to the LSST cadence and DR1 (Plan "A"), estimators that meet the criteria defined in § 4 and that have been fully validated during the commissioning period will be prioritized. If, however, the commissioning SV period is cut short, an *Early Science Period* of $\approx 3-6$ months that is different to regular survey operations may be enacted (Plan "B"). During this period, it is envisaged that estimators that meet the criteria defined in § 4 could continue to work on validation during this *Early Science Period*, and that a decision would be taken on which estimator(s) to run for DR1 prior to reverting to the LSST cadence and DR1. If a further shakedown of operations procedures and data taking is still required at handover to Operations (Plan "C"), the implications for photo-*z* estimators will be reviewed at that time based on the actual state of the system to determine the best course of action for photo-*z* estimators.

In any of these scenarios, estimators that will return the most accurate photo-*z* as early in the survey as possible and meet the criteria in (§ 3) may be prioritized. In the first year of LSST, it might be simpler to use a template-fitting photo-*z* estimator and avoid potential is-

sues related to computation resources and/or the need to train a machine learning model. Additionally, the likelihood that the large spectroscopic training sets needed for ML photo-z estimators exist will continue to increase through the 2020s. However, if a machine learning estimator is applied for LSST DR1 and DR2, it should be a community-accepted estimator with demonstrated success in other surveys, preferably surveys that overlap the LSST footprint, as this will facilitate the characterization and validation of the LSST photo-z.

C.5 A Summary of Scientific Recommendations from the LORs

As described in § 5.1, of the 6 LOR describing scientific use cases:

- 1 focused on galaxies
- 2 focused on dark energy
- 3 focused on active galactic nuclei

Recommendations that are within scope for the DM team (or were already in the plans or evaluation criteria; § 4) and are likely to be achieved:

- adopt a PZ estimator that provides PDFs, point estimates, SED/template fits, physical params, and flags
- adopt at least one PZ estimator that can generate accurate PZ for galaxies with AGN
- adopt a PZ estimator that performs well with imperfect priors / incomplete training sets (especially at low-z)
- provide PZ in the Object catalog at the time of data release (i.e., no delay)
- make all PZ-related code, inputs, metadata, etc. available (for reproducibility)
- correct for MW dust before passing fluxes to the PZ estimator
- provide users with a framework and tools to cross-match to non-LSST catalogs

Any of these recommendations which were not already part of the evaluation criteria in § 4 have been added.

Recommendations that present a challenge:

- incorporate non-LSST, non-optical photometry into the PZ estimates
- cross-match the LSST Object catalogs with external data sets (e.g., an X-ray AGN catalog) and provide the results with flags and identifiers
- release revised PZ data products every few months as new sky regions are observed
- link Object identifiers between data releases (so users can get new host data for old alerts)
- in the alert packets, provide host galaxy Object data (PZ, photometry, shape) and non-LSST survey data for host galaxies (spec-z, PZ, photometry)

These challenges are out of scope for the Rubin Construction project (and some violate the current Rubin data rights policy), and some items (such as providing cross-matched catalogs) are best left to the expertise of the science community. However, just because an idea or data product presents a challenge, or is beyond the scope of Rubin construction and commissioning, does not mean it could never be accomplished. Adopting and federating value-added user-generated data products, and updating algorithms in the LSST Science Pipelines, is an Operations-era (2024 and beyond) activity that will take input from the science community via the Rubin Users Committee, the Community Engagement Team, the Science Advisory Committee, and other such groups. It is very useful to have these suggestions raised early on so that the DM team can help the science community prepare for Operations; all of the LOR providing these recommendations are much appreciated.

D Original Call for Photo-z Letters of Recommendation

Deadline: Sep 30 2021.

Submission: In the spirit of open cooperation and collaboration towards our collective goal of generating LSST Object PZ that serve a wide variety of science goals, all letters will be public. Letters should be submitted by creating a new topic in the “Science – Photometric Redshifts” category at [Community.lsst.org](https://community.lsst.org)⁸. Use a title that starts with “LOR” and provides a bit of detail,

⁸Start at this link, <https://community.lsst.org/c/sci/photoz>, click “+New Topic” at upper right, enter a title and text (do not worry about tags), and if applicable upload a PDF using the up-arrow icon in the menu bar.

e.g., “LOR for the CMNN PZ Estimator”, “LOR: PZ and Local Volume Science”. Short letters could be presented in the text body of the new topic; longer letters could be provided by uploading an accompanying PDF.

Introduction: The Rubin Observatory Data Management (DM) team is tasked with constructing LSST Science Pipelines that produce science-ready data products, and this includes photometric redshift (PZ) estimates for the LSST data release (DR) Object catalog (e.g., `1s.st/dpdd`). As described in the LSST PZ Roadmap (`1s.st/dmtn-049`), DM will select one or more existing, community-vetted algorithms to generate Object PZ which, at least initially, meet a set of minimum scientific attributes and serve the widest variety of science applications.

LOR Purpose: The Rubin science community has a considerable wealth of expertise in generating PZ catalogs and will be the primary users of the Object PZ data products. These letters provide a formal opportunity for the science community to advocate for one or more photo-z estimators that will meet their minimum scientific needs, and/or to define these minimum needs so that DM can consider them when shortlisting PZ estimators.

LOR Writers: Any group or individual who would use the LSST Object PZ for their future scientific analyses, and/or are developers of potentially suitable PZ estimators, are encouraged to submit an letter. The DM team is especially interested in hearing the needs of scientists who plan to use the LSST Object catalog but who would/could not generate custom PZ estimates. LORs are not restricted to Rubin data rights holders.

LOR Scope: Letters could qualitatively (or quantitatively) recommend one or more specific PZ estimators (or type of estimator), but they do not have to: letters could instead recommend minimum attributes of the LSST PZ data product that would enable basic LSST science, or focus on describing the science that the LSST PZ data product should enable. Discussions on the research and development of new or improved algorithms can be considered as beyond the scope of these letters, as the shortlist will only include currently existing PZ estimators. New in-depth quantitative analyses of PZ estimator performance can also be considered as beyond the scope of the LORs because such analyses are the focus of the next stage of the PZ roadmap (see below). That said, no letters will be rejected for extending beyond these scope boundaries, which are provided just to guide and inform the community's efforts.

Guidelines for PZ LORs:

- Keep scope in mind: to identify the minimum scientific attributes, a wide set of science applications, and established PZ estimators for the LSST Object PZ data products.
- Follow the template. It's okay to skip some sections, but do not add new ones.
- Make it short (1-3 pages ideally) – a lot of detail is not needed at this time.
- Be qualitative – quantitative analyses will be the focus of the PZ Validation Cooperative.
- Refer to DMTN-049 for more detail about the roadmap, evaluation criteria, and LOR.

Beyond the LORs: The Rubin DM team will use these letters to inform the evaluation criteria used to select PZ estimators, and to assemble a shortlist of viable PZ estimators to be evaluated quantitatively using Rubin Commissioning data during the “PZ Validation Cooperative” phase of the PZ Roadmap. DM might add viable PZ estimators to the shortlist if DM thinks that they meet the evaluation criteria and the community’s scientific needs, even if they were not mentioned by any of the LORs. Writing a letter is not required (and will not be construed as a commitment) to participate in the PZ Validation Cooperative.

D.1 LOR Template

Title: E.g., “LOR on behalf of X science”, or “LOR for the X PZ Estimator”.

Contributors: Names and affiliations of the letter writers.

Co-signers: If applicable, the names and affiliations of co-signers.

1. Summary Statement – *Provide a short statement that introduces the writers’ and the letter’s main recommendation(s). Include the citations for any software discussed.*
2. Scientific Utility – Rubin DM seeks to understand LSST PZ-related science cases in order to ensure that the LSST Object PZ data products will be scientifically useful for a wide variety of communities, especially those which would/could not generate custom PZ. *Describe how you would use the PZ data products for your LSST science, or the LSST science enabled by the PZ estimator(s) you are recommending. Examples from past experiences would be useful here.*
3. Outputs – Rubin DM seeks to ensure that the selected LSST Object PZ estimator generates science-ready outputs that serve a wide variety of the community’s minimum scientific needs. *If possible, describe the minimum set of PZ outputs that are required for*

your science, or the outputs generated by the PZ estimator(s) you are recommending. E.g., full posteriors, point estimates, statistics (mode, mean, standard deviation, skewness, kurtosis), best-fit templates, and/or flags (e.g., quality, failure modes).

4. Performance – Rubin DM seeks a general understanding of the minimum quality of PZ estimates needed to meet the basic PZ-related science goals of the community. *If possible, describe the minimum PZ quality that would enable your LSST science (e.g., the minimum point-estimate error at intermediate redshifts, or whatever is relevant to your science goals), or the predicted minimum quality of the PZ estimator(s) you are recommending. It is understood that this information might not be available at this time.*
5. Technical Aspects – Rubin DM has a set of technical considerations for PZ estimators, regarding their scalability, inputs, outputs, language, external data sets, training, storage, and compute resources. *If possible – and this is probably only possible for letters being prepared by PZ algorithm developers – please briefly address the technical considerations described in Section 3 of the PZ Roadmap (ls.st/dmtn-049). Details are not necessary, but an evaluation of the technical considerations at the level of “will meet”, “will probably meet”, “probably will not meet”, “will not meet” would be most helpful at this time.*