



Tuning the Legacy Survey of Space and Time Observing Strategy for Solar System Science: Incremental Templates in Year 1

James E. Robinson¹ , Megan E. Schwamb² , R. Lynne Jones^{3,4} , Mario Jurić⁵ , Peter Yoachim⁵ , Bryce T. Bolin⁶ , Colin O. Chandler^{5,7,8,9} , Steven R. Chesley¹⁰ , Grigori Fedorets^{11,12} , Wesley C. Fraser¹³ , Sarah Greenstreet^{5,14} ,

Henry H. Hsieh¹⁵ , Lauren J. McGinley² , Stephanie R. Merritt² , Cyrielle Opitom¹ , and John K. Parejko¹⁶

¹ Institute for Astronomy, The University of Edinburgh, Royal Observatory Edinburgh, Blackford Hill, Edinburgh, EH9 3HJ, UK; james.robinson@ed.ac.uk

² Astrophysics Research Centre, School of Mathematics and Physics, Queen's University Belfast, Belfast BT7 1NN, UK

³ Rubin Observatory, 950 N. Cherry Ave., Tucson, AZ 85719, USA

⁴ Aerotek, Suite 150, 4321 Still Creek Drive, Burnaby, BC V5C 6S7, Canada

⁵ Department of Astronomy & the DiRAC Institute, University of Washington, 3910 15th Ave. NE, Seattle, WA 98195, USA

⁶ Eureka Scientific, Inc. 2452 Delmer Street, Suite 100, Oakland, CA 94602-3017, USA

⁷ LSST Interdisciplinary Network for Collaboration and Computing, 933 N. Cherry Ave., Tucson, AZ 85721, USA

⁸ Department of Astronomy & Planetary Science, Northern Arizona University, P.O. Box 6010, Flagstaff, AZ 86011, USA

⁹ Raw Data Speaks Initiative, USA

¹⁰ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109, USA

¹¹ Finnish Centre for Astronomy with ESO, University of Turku, FI-20014 Turku, Finland

¹² Department of Physics, University of Helsinki, P.O. Box 64, 00014 Helsinki, Finland

¹³ National Research Council, Herzberg Astronomy and Astrophysics Research Centre, 5071 W. Saanich Rd., Victoria, BC V9E 2E7, Canada

¹⁴ Rubin Observatory/NSF NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, USA

¹⁵ Planetary Science Institute, 1700 East Fort Lowell Rd., Suite 106, Tucson, AZ 85719, USA

¹⁶ Department of Astronomy, University of Washington, 3910 15th Ave. NE, Seattle, WA 98195, USA

Received 2024 November 25; revised 2025 February 26; accepted 2025 March 3; published 2025 June 20

Abstract

The Vera C. Rubin Observatory is due to commence the 10 yr Legacy Survey of Space and Time (LSST) at the end of 2025. To detect transient/variable sources and identify solar system objects (SSOs), the processing pipelines require templates of the static sky to perform difference imaging. During the first year of the LSST, templates must be generated as the survey progresses; otherwise, SSOs cannot be discovered nightly. The incremental template generation strategy has not been finalized; therefore, we use the Metric Analysis Framework (MAF) and a simulation of the survey cadence (`one_snap_v4.0_10yrs`) to explore template generation in Year 1. We have assessed the effects of generating templates over timescales of days–weeks, when at least four images of sufficient quality are available for $\geq 90\%$ of the visit. We predict that SSO discoveries will begin ~ 2 – 3 months after the start of the survey. We find that the ability of the LSST to discover SSOs in real time is reduced in Year 1. This is especially true for detections in areas of the sky that receive fewer visits, such as the North Ecliptic Spur (NES), and in less commonly used filters, such as the u and g bands. The lack of templates in the NES dominates the loss of real-time SSO discoveries; across the whole sky the MAF main-belt asteroid (MBA) discovery metric decreases by up to 63% compared to the baseline observing strategy, whereas the metric decreases by up to 79% for MBAs in the NES alone.

Unified Astronomy Thesaurus concepts: [Solar system astronomy \(1529\)](#); [Small Solar System bodies \(1469\)](#); [Surveys \(1671\)](#)

Materials only available in the [online version of record](#): [animations](#), [figure sets](#), [machine-readable tables](#)

1. Introduction

The planned observing strategy for the Vera C. Rubin Observatory's Legacy Survey of Space and Time (LSST) has been revised and optimized over the past 7 yr (Ž. Ivezić & The SCOC 2021; F. Bianco & The SCOC 2022, 2024; F. B. Bianco et al. 2022). With an expected start date in late 2025, the survey will span 10 yr, covering $\sim 18,000$ deg² in six broadband filters (u , g , r , i , z , and y) using the Rubin Observatory LSST Camera (LSSTCam), with a 9.6 deg² field of view (FOV) and the 8.36 m Simonyi Survey Telescope. The details of the survey's main science drivers and science requirements are summarized in LSST Science Collaboration et al. (2009), Ž. Ivezić & The

LSST Science Collaboration (2013), Ž. Ivezić et al. (2019), F. B. Bianco et al. (2022), and references within. The planned LSST observing strategy has primarily been assessed via simulations of the observing strategy, in order to determine the expected outcomes over 10 yr of observations. These simulations are generated by the Rubin operations simulator (`rubin_sim`) and the Rubin Observatory scheduler (`rubin_scheduler`), which combines the expected properties and performance of the telescope and camera with the conditions at the observing site using archival weather models for Cerro Pachón (A. J. Connolly et al. 2014; F. Delgado et al. 2014; LSST Science Collaboration et al. 2017; E. Naghib et al. 2019; R. L. Jones et al. 2020). The Rubin Metric Analysis Framework (MAF; R. L. Jones et al. 2014) takes a `rubin_sim` simulation as input and can be used to compute key values for a wide range of science cases, which then can be used to compare the performance of various LSST observing strategies.

Previous analyses of the LSST observing strategy (e.g., LSST Collaboration et al. 2017; R. L. Jones et al. 2018; M. Lochner et al. 2018; I. Andreoni et al. 2022; J. E. Gizis et al. 2022; E. D. Feigelson et al. 2023; M. E. Schwamb et al. 2023) have primarily focused on the scientific outcomes based on all data accumulated over the total 10 yr duration of the survey. In other words, they only consider the cumulative results obtained throughout the survey from the nightly stream of alerts, daily Rubin Prompt Processing (RPP) Data Products, or yearly data releases (described in M. Jurić et al. 2023). When making major decisions about how LSST will be executed (e.g., survey footprint area, distribution of observations between the optical filters, and observing cadence across different sky regions), computing and comparing MAF metrics (e.g., the total number of supernovae, main-belt asteroids, hereafter MBAs, or galaxies discovered) on 10 yr of observations in a given `rubin_sim` cadence simulation is sufficient. This is appropriate because for the majority of the survey, i.e., Year 2 onwards, it is expected that the generation of data products will be well described by the survey simulations described above. The main difference occurs in Year 1 of Rubin science operations when the survey will cover the sky for the first time and, as such, will not have previous LSSTCam images, which are required to describe the static background sky. This will impact the real-time discovery and detection of astrophysical transients; e.g., supernovae, variable stars, and moving objects. A further detailed analysis is needed to understand what will be discovered during Year 1 of LSST.

The nightly RPP pipelines will use difference imaging to find transient sources within each LSST observation, otherwise known as a visit. In difference imaging, a template representing the static sky, made through the combination of several past exposures, will be subtracted from an observation, thus creating an image of sources that represent changes relative to the static sky. Within 60 s of an image readout from LSSTCam, the resulting transient detections will be sent out as world public alerts in the Rubin alert stream (Ž. Ivezić & The LSST Science Collaboration 2013; Ž. Ivezić et al. 2019; E. Bellm et al. 2020; L. P. Guy et al. 2023; M. Jurić et al. 2023). The transient detections will also be fed into the Rubin Solar System Processing (SSP) pipelines in order to discover/recover moving solar system objects (SSOs), and report those detections to the Minor Planet Center (MPC; J. Myers et al. 2013; Ž. Ivezić et al. 2019; M. Jurić et al. 2020). In addition to nightly alerts, there will be annual Rubin Observatory data releases, where template images are regenerated, and all the survey images are reprocessed by the RPPs pipelines for transient detections. It is expected that each data release will have higher-quality template images for difference imaging, as the number of images available for templates increases with time.

This strategy will not work for Year 1 of LSST when the images required to produce templates are not yet available. A different approach is needed if astrophysical transients and SSOs are going to be found in near real time during the first year. To this end, the Rubin Data Management System will produce incremental templates as Year 1 of LSST progresses (L. P. Guy et al. 2023). The specific requirements, such as the number of observations required and how frequently the relevant data management pipelines will run to generate these Year 1 templates, are still to be decided (F. Bianco & The SCOC 2024). The selected strategy will have significant

implications for which of the Year 1 images can undergo difference imaging in the RPP pipeline (M. L. Graham et al. 2020; L. P. Guy et al. 2023). In particular, the Rubin SSP pipelines will only be able to search for moving objects on the night the observations are taken if a preexisting template is available. The detections will not be lost per se, as the Rubin data management pipelines will perform a complete reanalysis of all previous data when creating each annual data release. But this will limit the opportunities for immediate time-sensitive follow-up that would have been enabled if the SSOs discoveries were announced and preliminary orbits were available in the MPC. M. E. Schwamb et al. (2021) highlights some of the high-impact solar system science enabled by Year 1 LSST incremental template generation, such as the follow-up with ground- and space-based telescopes of interstellar objects (ISOs) passing through the solar system or potentially hazardous asteroids (PHAs) for planetary defense, in addition to catching cometary outbursts and temporary Earth satellites (known as “minimoons,” B. T. Bolin et al. 2020; G. Fedorets et al. 2020). M. E. Schwamb et al. (2023) examined how the LSST observing strategy parameters can be tuned to improve or benefit solar system science over the 10 yr of the survey and demonstrated the importance of evaluating various strategies for producing templates in Year 1 of Rubin operations. In recent years, the Rubin Observatory Survey Cadence Optimization Committee (SCOC) has come to a consensus for the majority of significant observing strategy decisions (Ž. Ivezić & The SCOC 2021; F. Bianco & The SCOC 2022, 2024). Now, we can begin to focus on the expectations for the discovery of SSOs during the first year of the survey and start exploring the trade-offs and benefits of potentially boosting incremental template production in Year 1 of LSST.

In this work, we investigate the generation of templates during Year 1 of LSST and assess the impact on SSO discovery metrics compared to the nominal LSST cadence simulation MAF metrics (where the existence of templates is assumed). In Section 2, we present our methodology for analyzing the `one_snap_v4.0_10yrs` Rubin cadence simulation, in order to identify which visits had an available template image at the time of observation. This allows us to create a list of visits in Year 1 that could have been processed by RPP to produce transient alerts, in particular detections of SSOs. We present our results of this analysis in Section 3 for several different timescales for incremental template generation and summarize our conclusions in Section 4.

2. Methods

2.1. Simulated LSST Pointings

In order to investigate the effects of incremental template generation, we make use of the Rubin MAF (R. L. Jones et al. 2014) and a simulated LSST survey pointing history generated with the `rubin_sim` operation simulator. This uses the `rubin_scheduler`, with inputs for the predicted performance of the telescope, optics, and LSSTCam, alongside site weather models and values for Cerro Pachón’s sky brightness (P. Yoachim et al. 2016). In addition, the simulation includes realistic periods of downtime for maintenance and engineering work. The `rubin_sim` operation simulator and `rubin_scheduler` are presented and reviewed in detail in A. J. Connolly et al. (2014), F. Delgado et al. (2014), F. Delgado & M. A. Reuter (2016), P. Yoachim et al. (2016),

LSST Collaboration et al. (2017), R. L. Jones et al. (2018), E. Naghib et al. (2019), R. L. Jones et al. (2020), F. B. Bianco et al. (2022), and references therein. We chose to use the output from the `one_snap_v4.0_10yrs` cadence simulation (F. Bianco & The SCOC 2024; P. Yoachim 2024) in our analysis. At the time of submission (2024 November), this simulation was the most accurate representation of how LSST will be conducted, having been refined by the Rubin SCOC in consultation with the various LSST science collaborations and Rubin data rights community. The details of the earlier `baseline_v3.0_10yrs` are laid out in F. Bianco & The SCOC (2022), where a number of minor adjustments have been incorporated into the survey strategy from previous versions to produce `one_snap_v4.0_10yrs` as described in F. Bianco & The SCOC (2024). In particular, `one_snap_v4.0_10yrs` simulates running the survey with a single exposure or “snap” for each visit, whereas previous simulations had assumed that each visit would be divided into two snaps that would be later coadded by the Rubin data management pipelines. A single snap reduces the overhead due to an image readout, increasing the time available to the survey by $\sim 7\%–9\%$ (see Section 3.6 of F. Bianco & The SCOC 2024). The SCOC has already endorsed switching to the one snap per visit observing mode if it is technically feasible after commissioning tests with LSSTCam. The main reason to have two snaps per visit was for aiding in cosmic-ray identification and rejection. Today, there are several automated methods that can identify cosmic-ray strikes in single charged-coupled device (CCD) images (e.g., J. E. Rhoads 2000; P. G. van Dokkum 2001; L. Shamir 2005; C. McCully et al. 2018). Thus, we expect that the one snap cadence simulation will best represent the survey observing strategy during the LSST’s first year on the sky.

Full details on the observing cadence of the `one_snap_v4.0_10yrs` simulation is described in F. Bianco & The SCOC (2024). We highlight a few of the notable aspects of the observing strategy and changes compared to previous `rubin_sim` simulations. The `one_snap_v4.0_10yrs` simulation includes more accurate estimates for the throughput of the optical system with silver instead of the originally planned aluminum mirror coatings. All filters except for u benefit from the switch to silver. Because of the revised throughputs, g , r , i , z , and y -band exposures are 29.2 s whereas u exposures were lengthened to 38 s. More visits were shifted to the u filter; the number of u -band exposures was increased by 10% compared to `baseline_v3.0_10yrs` simulation. We note that the gain in observing time from switching to one snap is not available in Year 1 of the `one_snap_v4.0` simulation, as 8 weeks worth of downtime was included to account for additional engineering time that is expected to be needed postcommissioning (F. Bianco & The SCOC 2024). This engineering time is more heavily weighted to the start of the survey and tapers off toward the end of Year 1. There is no rolling cadence (where the sky is divided into bands, and during “on” years, the band gets a more concentrated number of visits) implemented in the first year of `one_snap_v4.0` due to the need for generating templates across the entire survey footprint for Year 2 operations. The distributions of on-sky visits in Year 1 are shown in Figure 1.

The LSST is comprised of multiple surveys (see Figure 1) with the largest component being the $\sim 18,000$ deg² Wide-Fast-Deep (WFD); see Ž. Ivezić & The SCOC (2021),

F. B. Bianco et al. (2022), and F. Bianco & The SCOC (2022) for a detailed description. In this analysis, we have excluded any visits associated with the five Deep Drilling Fields (DDFs) or the low-solar elongation twilight survey.¹⁷ These visits could skew the MAF statistics because these surveys within the LSST footprint follow observational strategies that differ significantly from the observing strategy of the main WFD survey and the other additional smaller surveys making up LSST; the North Ecliptic Spur (NES) and dusty (Galactic) plane, and south celestial pole minisurveys. The NES minisurvey observes the sky northward of the WFD footprint, up to $+10^\circ$ of the ecliptic in the northern hemisphere sky, in $griz$ filters. We note that the NES is observed with fewer total visits than the WFD observing strategy. M. E. Schwamb et al. (2018b), M. E. Schwamb et al. (2023), and references therein provide a detailed description of the NES and its importance for LSST solar system science. The DDFs have dense temporal coverage over a very localized portion of the sky (one camera pointing, or two in the case of the Euclid DDF), and so, any SSOs in a DDF would receive many more detections than is representative for the vast majority of SSOs imaged in LSST. Similarly, the twilight observations are taken at low-solar elongation in order to find small bodies on orbits interior to the Earth, which translates to camera pointings that are highly constrained on sky and temporally restricted to when the Sun is between 0° and -12° elevation (F. Bianco & The SCOC 2024).

From this point on, we shall refer to the Year 1 baseline as `one_snap_v4.0`, which is the `one_snap_v4.0_10yrs` pointing history with the DDF visits, low-solar elongation twilight survey, and visits after Year 1 removed. We only consider visits during Year 1 of the survey as after this time an interim incremental template strategy will no longer be necessary; full sky coverage with sufficient quality for template generation is expected to be achieved after Year 1 (L. P. Guy et al. 2023). We take the modified `one_snap_v4.0` simulated pointing history and count the number of visits suitable for generating templates matching our requirements described in Section 2.3. We assume that the RPP and SSP pipelines will only generate alerts and solar system detections for areas of the sky with templates, as difference imaging cannot be performed without templates (M. L. Graham et al. 2020). Figure 2 shows the per filter `one_snap_v4.0` simulation’s Year 1 visit sky maps that are considered for our analysis of incremental template generation timescales. Figure 3 displays a snapshot from an animation showing the nominal sky coverage of r -band observations expected over the first year of the survey.

2.2. Tracts and Patches

Each night, the Rubin scheduler will randomly select a goal LSSTCam rotator angle between -80° and $+80^\circ$ for the entire night of observations (P. Yoachim 2024). The scheduler will try to execute observations with position angles as close as possible to the desired angle. Visits to the same pointing within a night will have close but slightly different camera rotator positions. The scheduler will also execute spatial camera dithers between subsequent visits to the same on-sky LSST

¹⁷ Structured Query Language query for the `one_snap_v4.0_10yrs` pointing database: `select * from observations where night < 365 and scheduler_note not like "%DD%" and scheduler_note not like "%twilight%"`.

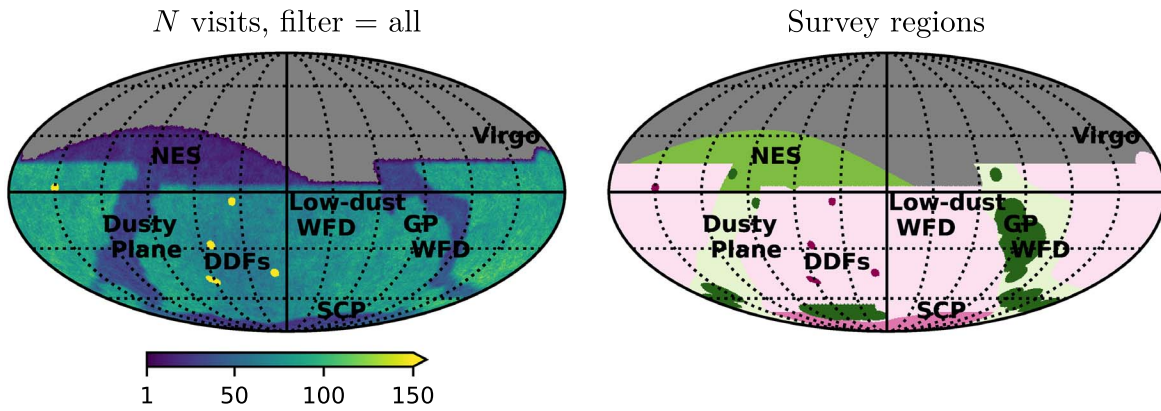


Figure 1. Left: A sky map showing the total number of visits, in all filters, at the end of Year 1 for the `one_snap_v4.0_10yrs` observing strategy (Mollweide projection). This sky map was generated using a HEALPIX (Hierarchical Equal Area isoLatitude Pixelization; K. M. Górski et al. 2005) resolution of `nside = 256`. The plots are centered on R.A. $\alpha = 0^\circ$ and decl. $\delta = 0^\circ$, with R.A. increasing to the left. R.A. and decl. lines are marked every 30° . The main sky regions of the survey are labeled as follows: Low-dust Wide-Fast-Deep (WFD), North Ecliptic Spur (NES), Galactic plane (GP) WFD, dusty (Galactic) plane, south celestial pole (SCP), Deep Drilling Fields (DDFs), and the Virgo cluster. Right: the color map denotes the on-sky extent of the different survey regions.

field (P. Yoachim 2024). These spatial and rotational dithers will lead to nonuniform coverage at subdetector length scales; therefore, the creation of a template image will not be as simple as selecting a number of images in a given filter at the same field pointing. Due to the shape of the LSSTCam footprint, the gaps between CCD detectors, and gaps between rafts (groupings of CCDs within LSSTCam), the dithering between observations requires that the LSST data management and reduction pipelines reduce noncoadded observations at the individual CCD level. Thus, the creation of image subtraction templates will also be performed on a size scale similar to LSSTCam CCDs. A similar strategy has been successfully employed in the Subaru Telescope’s Hyper Suprime-Cam Survey data reduction pipelines (J. Bosch et al. 2018).

The LSST data management pipelines divide the sky into a common pixel grid of overlapping square 1.6×1.6 tiles, dubbed “tracts,” which are themselves subdivided into 7×7 “patches” (see Figure 4; and J. Bosch et al. 2018; J. Swinbank et al. 2020). It takes nine tracts to fully cover a single LSSTCam pointing, which spans 3.5 deg^2 . A single patch is the approximate size of an LSSTCam CCD detector with dimensions of 13.7×13.7 . A patch is comprised of 4100×4100 pixels with a pixel scale of 0.2 per pixel, matching that of LSSTCam (Ž. Ivezić et al. 2019; J. Swinbank et al. 2020), and each patch overlaps by 100 pixels on a side with their neighboring patches. For a detailed overview of tracts and patches, we refer the reader to the summary paper for the Rubin Observatory’s Data Preview 0.2/Dark Energy Science Collaboration (DESC) Data Challenge 2 (DC2; LSST Dark Energy Science Collaboration 2021, hereafter LSST DESC).

A patch is the smallest unit that will be handled by Rubin data-processing pipeline. Template generation and, subsequently, image subtraction will be performed at the patch level. Figure 5 shows an example of the coverage in a single filter for a randomly selected patch chosen from the simulated Data Preview 0.2/DC2 (LSST DESC) Year 1 Data Release image templates. The impact from rotational dithers, spatial dithers, chip gaps, raft gaps, masked pixels at detector edges, and saturated sources can be seen. An LSST patch will therefore not have uniform coverage across all of its pixels. This must be accounted for when estimating the Year 1 incremental template production rates.

2.3. HEALPIX Sky Maps

To (1) track the Year 1 visits in a given area of the sky suitable for making an incremental template and (2) to identify which observations have sufficient template coverage to produce alerts and solar system detections in Year 1, we partition the sky using a HEALPIX map (Hierarchical Equal Area isoLatitude Pixelization¹⁸; K. M. Górski et al. 2005) as shown in Figures 1 and 2. This pixelization produces subdivisions of a spherical surface in which each healpixel has equal surface area and a resolution determined by `nside`, such that the whole sky is divided into $12 \times \text{nside}^2$ healpixels. It is too computationally expensive to consider incremental template generation at the individual patch pixel level as each patch has 16,810,000 pixels. Our best compromise is to instead focus on the patch as our smallest size element and use `nside = 256`. This results in healpixels with a resolution of approximately 13.7 , which is comparable in angular size to a patch. We note that this results in a HEALPIX sky map grid that is similar to, but not exactly aligned with, the patch/tract tessellation that the Rubin Observatory Data Management pipelines are using. By using a healpixel resolution comparable to the patch size, we on average balance the problems of oversampling (large `nside`, high resolution healpixels) and undersampling (small `nside`, low resolution healpixels) when making our HEALPIX incremental template coverage sky maps.

2.4. Year 1 Template Generation Timescales

As Year 1 of the survey progresses and images are taken according to the predefined survey strategy, the sky coverage increases nightly. However, for operational reasons (such as constraints on staffing and computational resources), template production is unlikely to occur nightly, but instead only on certain nights with some timescale (e.g., days to weeks). Template generation timescales have not yet been finalized by the Rubin Observatory Operations and Data Management Teams, but they are planning for a regular schedule (e.g., approximately monthly; M. L. Graham et al. 2020; L. P. Guy et al. 2023). We therefore explore a range of template generation timescales, $\Delta t = 3, 7, 14,$ and 28 day intervals from the start of the `one_snap_v4.0` simulation.

¹⁸ <http://healpix.sourceforge.net>

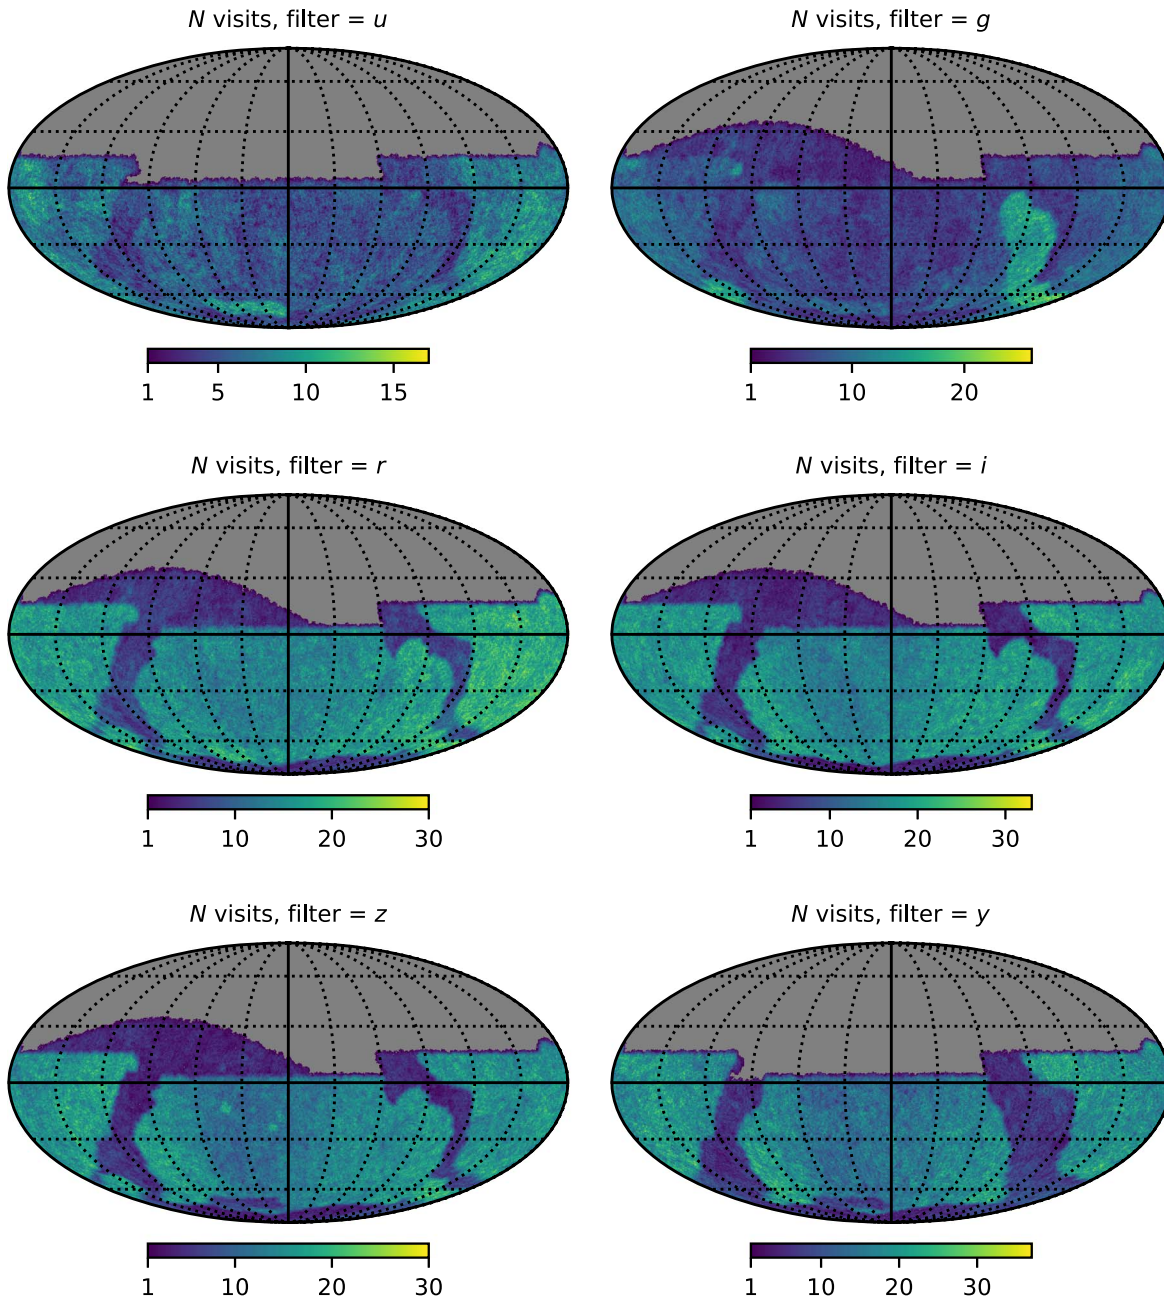


Figure 2. Sky maps of the total number of visits per filter accumulated during Year 1 of the `one_snap_v4.0` cadence simulation. Visits associated with DDFs and the low-solar elongation twilight survey have been removed. In each panel, we display the number of visits in each filter (*ugrizy*) to highlight the variations in footprint area and coverage (note the different ranges of each color bar). The sky maps were generated using a HEALPix resolution of $n_{\text{side}} = 256$.

2.5. Requirements for Incremental Template Production

The individual image requirements and total number of observations needed to make suitable Year 1 templates in each filter have not yet been finalized by the Rubin Observatory Data Management Team (M. L. Graham et al. 2020; L. P. Guy et al. 2023). On-sky tests with the commissioning camera (ComCam) and LSSTCam are needed. ComCam was installed in 2024 August, and the first photons were captured in 2024 October. At the time of this paper’s submission, LSSTCam on-sky commissioning is expected to start in spring 2025.¹⁹ In the meantime, we can make some reasonable assumptions based on

the expected performance of the telescope and camera system. There are benefits and trade-offs to balance for the image requirements applied to Year 1 template generation. This study is a first step in understanding the impact of incremental template generation, so we elect to use very broad requirements for image quality based on the simulated seeing and depth of each visit in `one_snap_v4.0`. Further work will be needed to explore the full phase space of image requirements to optimize LSST Year 1 template building.

2.5.1. Image Seeing and Depth Constraints

Not every LSST observation taken in the first year of science operations should be combined to produce an image subtraction template. Observations taken in poor seeing and/or bad

¹⁹ The latest Rubin Observatory construction milestone schedule can be found at <https://www.lsst.org/about/project-status>.

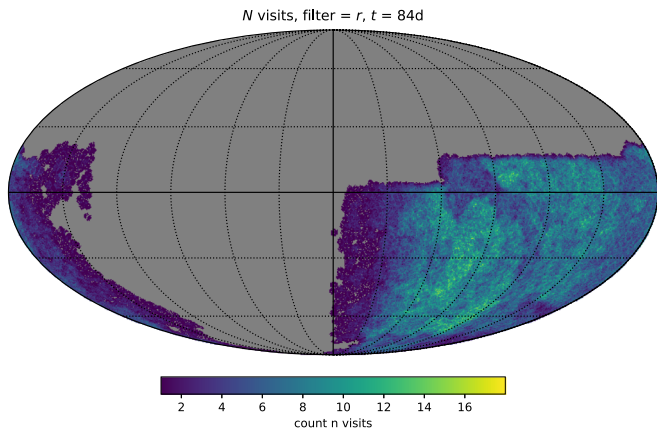


Figure 3. A snapshot from an animation showing the Year 1 sky coverage over time for the `one_snap_v4.0` observing strategy. The animation steps through the first year of the simulated LSST in intervals of 7 days, displaying the cumulative number of on-sky visits in the *r* filter. The plots are centered on R.A. = 0° and decl. = 0°. R.A. and decl. lines are marked every 30°. (An animation of this figure is available in the [online article](#).)




Figure 4. A cartoon schematic of tracts and patches in the LSST on-sky data management tessellation. The sky is divided into overlapping tracts each 1.6×1.6 . Each tract is comprised of 49 overlapping patches. Patches are roughly the size of an individual LSSTCam CCD detector with dimensions of 13.7×13.7 .

photometric conditions should ideally be excluded to minimize the number of artifacts in difference images. On the other hand, if image quality constraints are too restrictive, then very few observations will meet our requirements making templates even scarcer in Year 1. Using the values per visit reported in the `one_snap_v4.0` pointing database generated by `rubin_sim`, we evaluate the seeing quality using the “effective” full width at half-maximum (FWHM), which is the FWHM of a single Gaussian describing the point-spread function (PSF), and image depth from the 5σ limiting magnitude.²⁰ We have not directly included air-mass constraints on incremental template generation. We also set no minimum time threshold between suitable images as SSP only attempts to link sources that have moved within a single observing night. Most SSOs would be expected to move sufficiently far on the sky between the template building observations and future observations of the field to be identified as transient sources after the template is available. However,

²⁰ In `one_snap_v4.0`, these are the fields `seeingFwhmEff` and `fiveSigmaDepth`, respectively.




Figure 5. A representative patch randomly selected from the Rubin Data Preview (DP) 0.2/DESC DC2 (LSST DESC) simulated data products Year 1 annual templates. The extent of a single patch is 13.7×13.7 . The color bar reflects how many visits went into that patch pixel. The nonuniform coverage is a result of overlapping detector footprints from multiple visits taken with spatial and rotational dithers.

Table 1

Observation Seeing Statistics for the “Effective” FWHM Seeing: Median, Standard Deviation, Minimum, and Maximum

Filter	med (Seeing) (arcsec)	std (Seeing) (arcsec)	min (Seeing) (arcsec)	max (Seeing) (arcsec)
<i>u</i>	1.101	0.344	0.482	4.380
<i>g</i>	1.058	0.350	0.502	4.396
<i>r</i>	0.947	0.318	0.505	4.246
<i>i</i>	0.902	0.280	0.487	3.459
<i>z</i>	0.908	0.291	0.508	4.493
<i>y</i>	0.940	0.246	0.539	3.757

Note. Results are given per filter for Year 1 of the `one_snap_v4.0` cadence simulation (excluding low-solar elongation twilight survey and DDF visits). Standard deviation (std).

(This table is available in machine-readable form in the [online article](#).)

Table 2

Observation 5σ Limiting Magnitude (Image Depth) Statistics: Median, Standard Deviation, Minimum, and Maximum

Filter	med (5σ Depth) (mag)	std (5σ Depth) (mag)	min (5σ Depth) (mag)	max (5σ Depth) (mag)
<i>u</i>	23.497	0.353	20.523	24.515
<i>g</i>	24.461	0.549	21.704	25.589
<i>r</i>	24.168	0.423	21.833	25.054
<i>i</i>	23.658	0.390	21.667	24.617
<i>z</i>	23.079	0.348	20.421	23.894
<i>y</i>	22.076	0.308	20.240	22.873

Note. Results are given per filter for Year 1 of the `one_snap_v4.0` cadence simulation (excluding low-solar elongation twilight survey and DDF visits). Standard deviation (std).

(This table is available in machine-readable form in the [online article](#).)

Table 3
Solar System Object Discovery MAF Metrics Used in This Analysis

Population	Discovery Metrics	Metrics
PHAs		Three nightly pairs in 15 night discovery completeness for $H \leq 16.0$ Three nightly pairs in 15 night discovery completeness for $H \leq 22.0$
NEOs		Three nightly pairs in 15 night discovery completeness for $H \leq 16.0$ Three nightly pairs in 15 night discovery completeness for $H \leq 22.0$
MBAs		Three nightly pairs in 15 night discovery completeness for $H \leq 16.0$ Three nightly pairs in 15 night discovery completeness for $H \leq 21.0$
TNOs		Three nightly pairs in 15 night discovery completeness for $H \leq 6.0$ Three nightly pairs in 15 night discovery completeness for $H \leq 8.0$
OCC with $q \leq 5$ au		Three nightly pairs in 15 night discovery completeness for $H \leq 8.0$ Three nightly pairs in 15 night discovery completeness for $H \leq 17.0$
OCCs with $q \leq 20$ au		Three nightly pairs in 15 night discovery completeness for $H \leq 8.0$ Three nightly pairs in 15 night discovery completeness for $H \leq 12.0$

Note. These metrics are calculated for the “bright” and “faint” components of each population, defined by absolute magnitude H .

this may lead to contamination of the template images by slow moving sources, such as distant outer SSOs.

We impose broad image quality constraints to exclude the observations with the worst seeing and limiting magnitude. Visits of sufficient quality must all have their seeing within a given ratio, and the limiting magnitudes within a given range. Each filter is evaluated separately, as each filter has different expected image depth and seeing distributions, as shown in Tables 1 and 2. For a given healpix, filter f , and date in the survey, we look at all the relevant observations taken up to that point, and every observation j that matches the following criteria will be considered to build the Year 1 template at that healpixel location:

1. seeing of frame j /min (seeing of all available observations in filter f) < 2 ;
2. ($\max(5\sigma$ limiting magnitude of all available observations in filter f) $- 5\sigma$ limiting magnitude of frame j) < 0.5 .

This will reject the lowest-quality images on average. There is a chance a healpixel gets unlucky, and all the observations taken up to that date are poor quality. In that case, if the seeing and depths do not vary significantly between the exposures, a Year 1 template would still be produced for that healpix. We examine the quality of the templates produced in our simulation in Section 3.

2.5.2. Number of Suitable Observations Used to Build a Year 1 Template

We also need to decide the minimum number of observations meeting our quality assurance criteria (described in Section 2.5.1) that should be combined to produce the Year 1 incremental templates at the patch level. When considering healpixels with a comparable angular size to patches, we note that there will be some cases when the overlap of images that lies within a given healpixel will not fill 100% of the area (see the discussion in Section 2.2), and the counted number of visits suitable for generating a template will be slightly overestimated or underestimated. We expect that across many healpixels these effects will on average balance out. We can also help mitigate this effect by picking a reasonable threshold for the number of suitable observations that must overlap with a patch (our sky

map healpixel) to build a template such that at least 80%–90% of the patch is covered. Recent Rubin Observatory technical notes propose that three good quality observations might be sufficient for producing template images (M. L. Graham et al. 2020; L. P. Guy et al. 2023). The majority of the randomly selected DP0.2/DC2 patch in Figure 5 is covered by at least five observations with more than $\sim 80\%$ of the patch having coverage of at least four observations. We therefore add one more observation beyond the minimum proposed in M. L. Graham et al. (2020) and L. P. Guy et al. (2023) and require four Year 1 observations per patch per filter that meet or exceed our seeing and image depth thresholds to ensure successful template building.

2.6. Tracking Incremental Template Generation

We step through the Year 1 simulated rubin_simone_snap_v4.0 pointing database in Δt time intervals. At each template building session, n , we use MAF to iterate over all healpixels in our sky map and each of the six LSST filters, selecting all visits that span a given healpixel and separating these visits into possible template building or science observations. We define t_n as a date on which incremental templates are generated:

$$t_n = t_0 + \Delta t \cdot n \quad (1)$$

where t_0 is the survey start date. In this analysis, we investigated values of $\Delta t = 3, 7, 14,$ and 28 days (see Section 2.4). Templates must be created separately for each filter. As such, for each healpixel (aka patch), we select all overlapping visits in a given filter with the visit exposure time $t \leq t_n$. We assume that templates were last generated on the date t_{n-1} ; therefore, we divide all visits into possible template images ($t < t_{n-1}$) and science images ($t_{n-1} < t \leq t_n$). For the possible template images, we select only the visits that match our image quality criteria (see Section 2.5.1). If at least four images (see Section 2.5.2) fitting these criteria are overlapping the healpixel, then the template for that healpixel is assumed to have been generated at t_{n-1} .

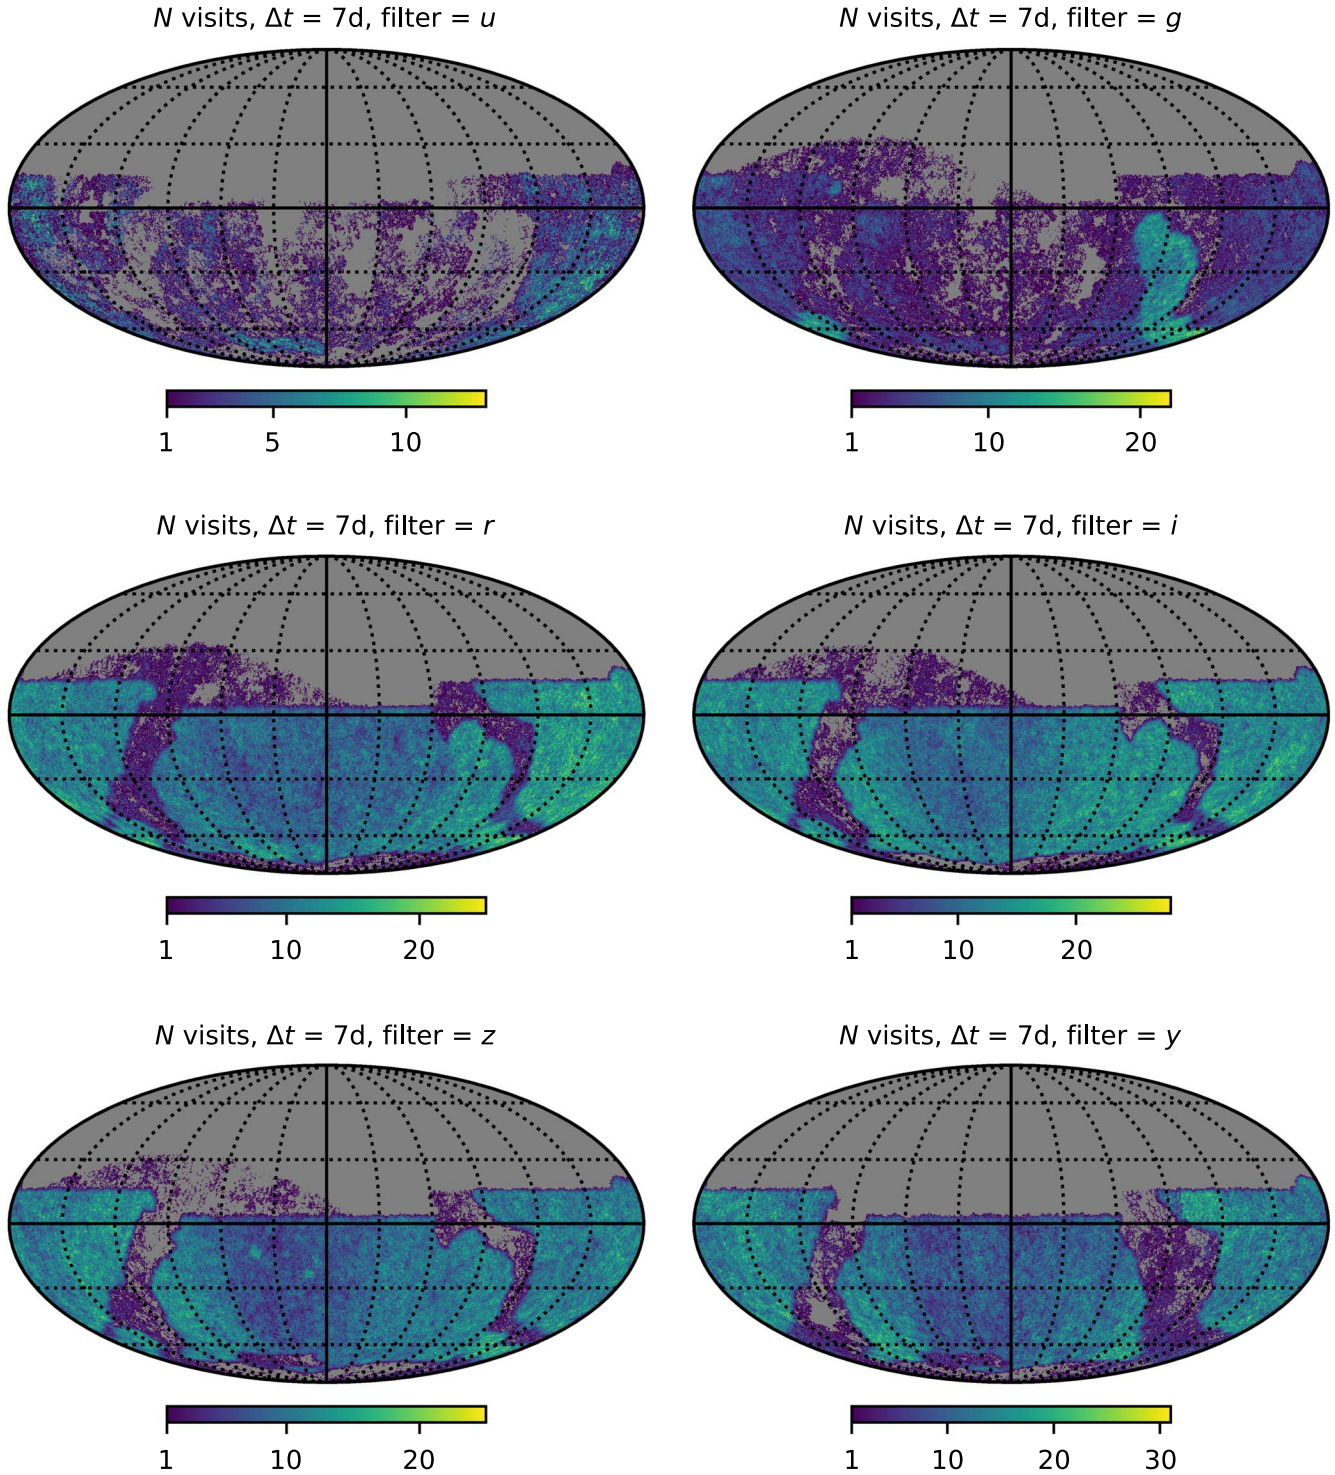


Figure 6. Sky maps showing the total number of visits with templates, per healpixel ($n_{\text{side}} = 256$), at the end of Year 1, assuming a template generation timescale of 7 days. Each panel shows the results for the *ugrizy* filters. Compare to Figure 2 for the baseline number of visits per pixel at the end of Year 1.

In this case, all subsequent visits to that healpixel ($t_{n-1} < t \leq t_n$) are assumed to have been successfully processed by the RPP and SSP pipelines and to have generated real-time transient detections, and therefore can be counted when assessing any metrics. If the template was not generated at t_{n-1} , then all subsequent visits to that healpixel are assumed to have not been processed, and they are discounted when assessing metrics. We note that some healpixels may have >4 suitable observations

available by the time of template generation and would likely have better depth/seeing and so be of higher quality. We do not adjust the 5σ limiting magnitude of an LSST observation based on the properties of its associated healpixels' templates. If the majority of the healpixel templates go to a much shallower depth than the observation it is subtracted from, there will be many bogus transient detections in addition to the real transient astrophysical sources and moving SSOs. In principle, the SSP

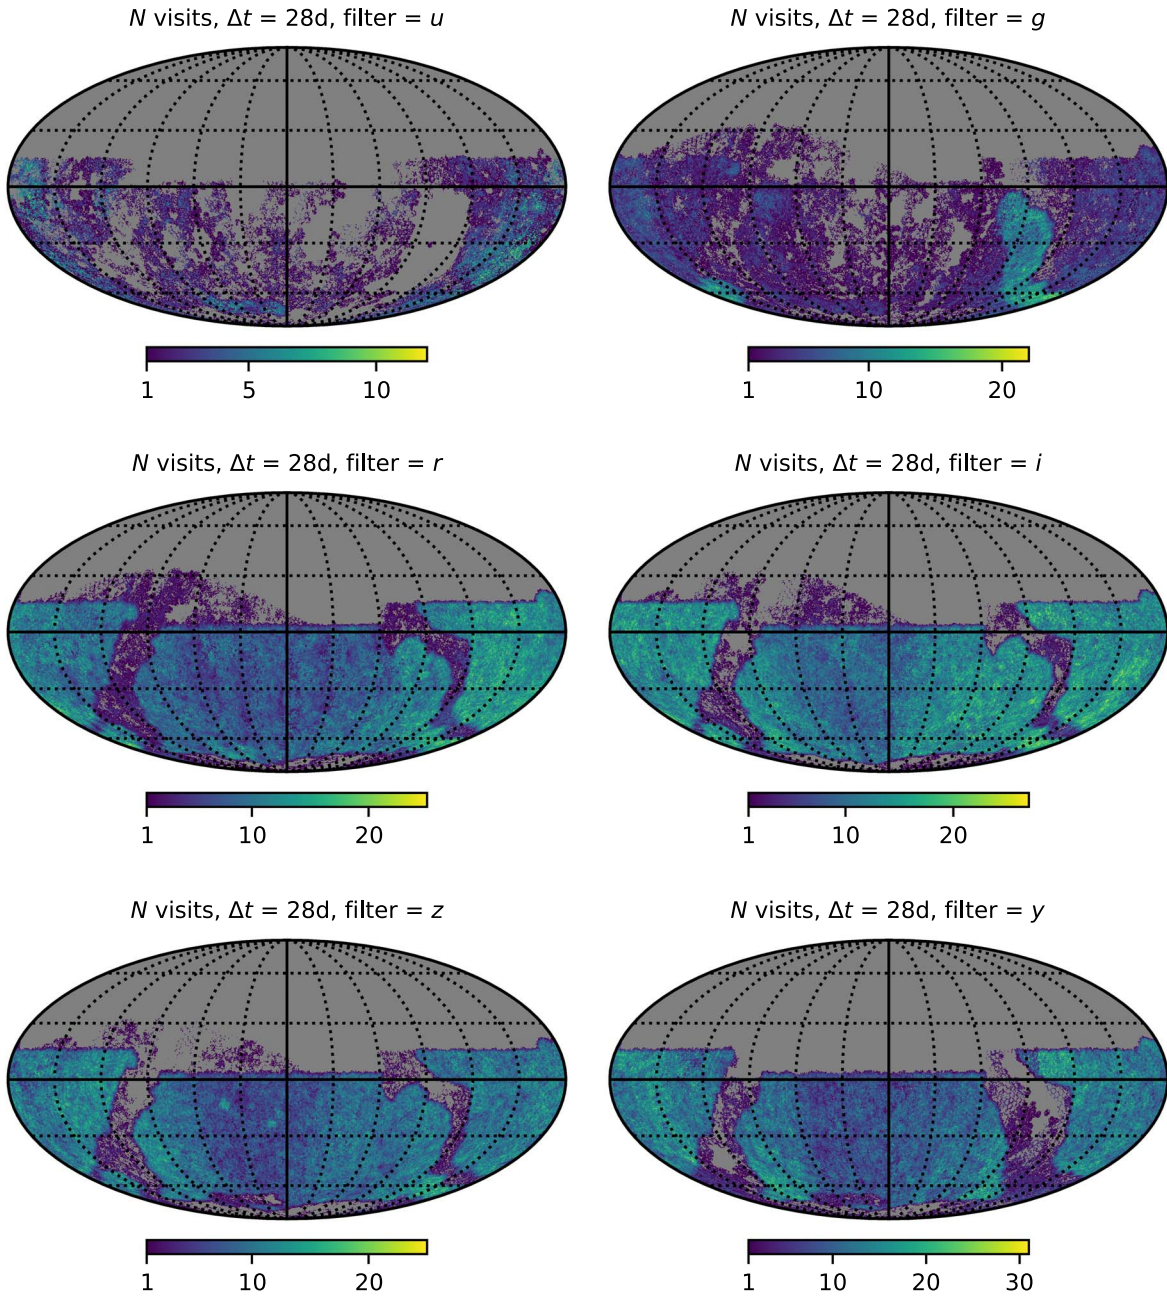


Figure 7. Same as Figure 6 for a template generation timescale of 28 days. Compare to Figure 2 for the baseline number of visits per pixel at the end of Year 1.

linking algorithms, namely, `HelioLinc3D`,²¹ should be able to match detections over a period of days to weeks and filter out those that are inconsistent with the Keplerian motion across the sky exhibited by a genuine SSO (M. J. Holman et al. 2018; A. Heinze et al. 2022). The number of random false-positive detections that could be combined in such a way to produce a decent heliocentric orbit fit should be quite low; however, we acknowledge that the presence of bogus detections could complicate the identification of genuine SSOs.

2.7. Applying MAF Metrics

At this stage, we are able to use MAF to run various metrics directly on the healpixels covering observations that had

templates available. We define the `reduceCount` metric as the total number of visits overlapping a given healpixel that meet certain criteria (such as filter) in a given time period. Given our methodology described above, we are assessing metrics on all observations up to the date t_n , where t_{n-1} was the last time templates were generated. Thus, with each time step in our simulation, we are effectively recording metrics for the periods of time $t_{n-1} < t < t_n$. As such, to get the final results for total number of visits per healpixel in Year 1, we must sum the values of the `reduceCount` metric, which were recorded at every t_n . We have also recorded the number of nights between template generation and the first visit to the healpixel, which we will refer to as `deltaNight`.

We also need to look beyond the individual healpixel (patch) level to the coverage at the LSSTCam FOV scale. At `nside = 256`, each visit contains ~ 183 healpixels. If only a

²¹ <https://github.com/lst-dm/heliolinc2>

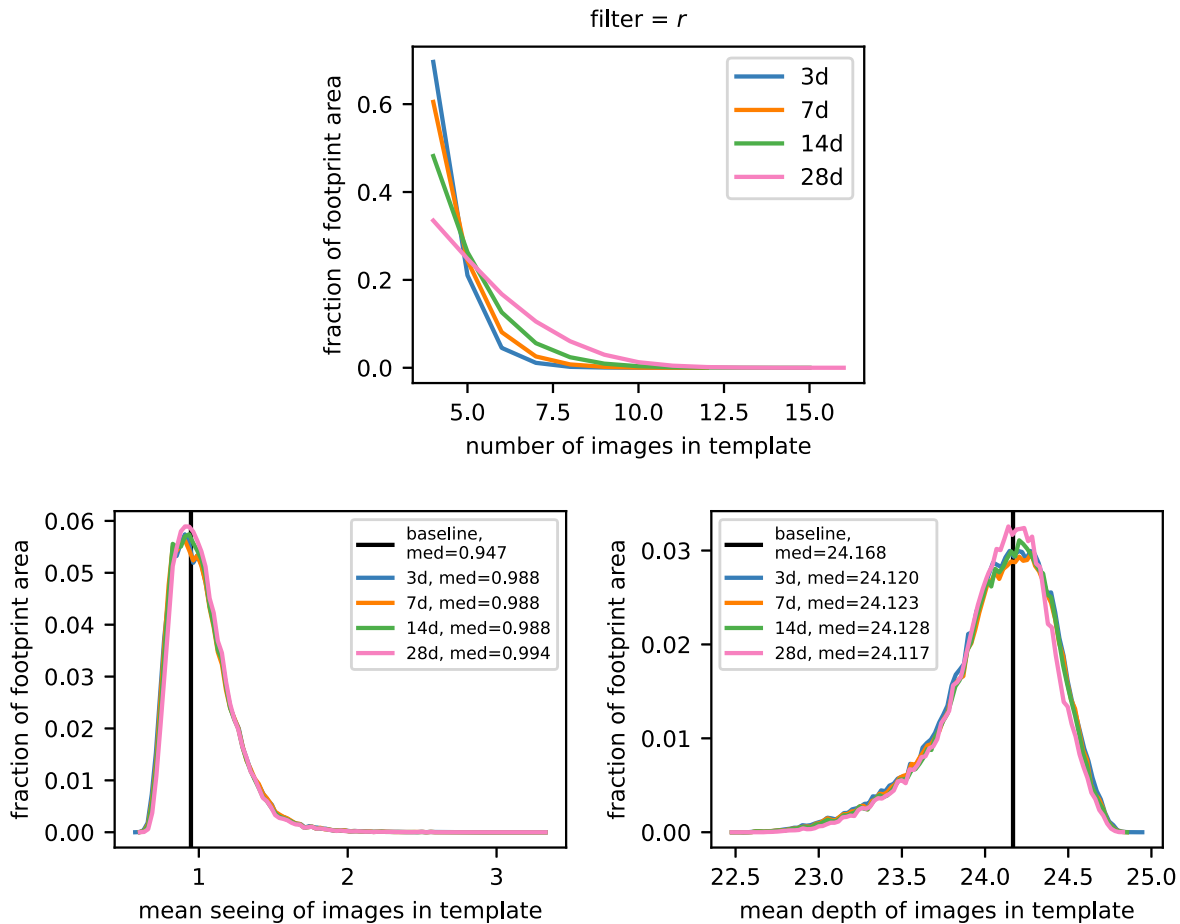


Figure 8. Histogram plots showing the quality statistics of the images incorporated into template images across the sky (as a fraction of unique footprint area in a given filter) for various generation timescales. Results are shown here for the r filter. The upper panel shows the number of images used in a template, where a minimum of four was required in these simulations. The lower left and lower right panels show the distribution of mean seeing and mean limiting magnitude (depth) of images used for templates, respectively. The solid vertical line indicates the median value of all images in the baseline survey (Tables 1 and 2). In these simulations, templates are considered on a healpixel basis, and, as such, the y-axis indicates the total sky area of template healpixels with a given image quality statistic value. (The complete figure set (6 images) is available in the [online article](#).)

few patches in a given LSST exposure are able to be image subtracted, then we expect moving object discoveries to be missed across most or all of the frame. This is because the SSP detection algorithm requires nightly tracklets for discovery; these are potential linkages from linear extrapolation using detections from two different observations of the same field in a single night (J. Myers et al. 2013; M. Jurić et al. 2020). SSP then takes the tracklets from the last 15 nights, and attempts to associate three tracklets onto a heliocentric orbit in order to identify new SSOs (J. Kubica et al. 2007; J. Myers et al. 2013; M. Jurić et al. 2020). As such, each visit must have reasonable template coverage to have good chances of detecting the moving object over the days–weeks required for SSP linkage. To this end, while simulating template generation as described above, MAF was used to count the number of constituent healpixels within each visit in the period $t_{n-1} < t < t_n$ that already had a template at t_{n-1} . From the number of template healpixels, we calculated the fractional template coverage of each visit (see Section 3.2) assuming an LSSTCam FOV of 9.6 deg^2 and healpixel area of $5.25 \times 10^{-2} \text{ deg}^2$ for $n_{\text{side}} = 256$. For discovering SSOs, we only run solar system MAF metrics on observations where the majority of the image area is available for template generation. For this study, we create a redacted `one_snap_v4.0` pointing database, where visits

with $<90\%$ template coverage have been removed, and use this to calculate a subset of MAF solar system metrics. We discuss the validity of selecting a threshold value of 90% template coverage in Section 3.2.

MAF is able to calculate a wide range of metrics related to solar system science, in order to evaluate the performance of the planned LSST observing cadence and the impact of tuning various survey strategy parameters in the Rubin scheduler. The detailed methodology for how the solar system metrics are calculated is described in M. E. Schwamb et al. (2023); we provide a brief overview here. Using the `movingObjects` module within MAF, a subsample of orbits are selected from a particular small body model population. Each orbit is then cloned for a range of absolute magnitude (H) bins, and broadband colors appropriate for the population are assigned. For each synthetic small body, ephemerides are generated and used to determine in which of the simulated LSST observations will the “object” be positioned within the camera footprint, if any. For those objects within an observation’s LSSTCam FOV, the apparent magnitude is calculated and combined with other relevant information about the visit stored in the `rubin_sim` simulated pointing database to determine whether the synthetic object would be bright enough for the RPP pipelines to detect. Using this data, MAF can calculate the fraction of successfully

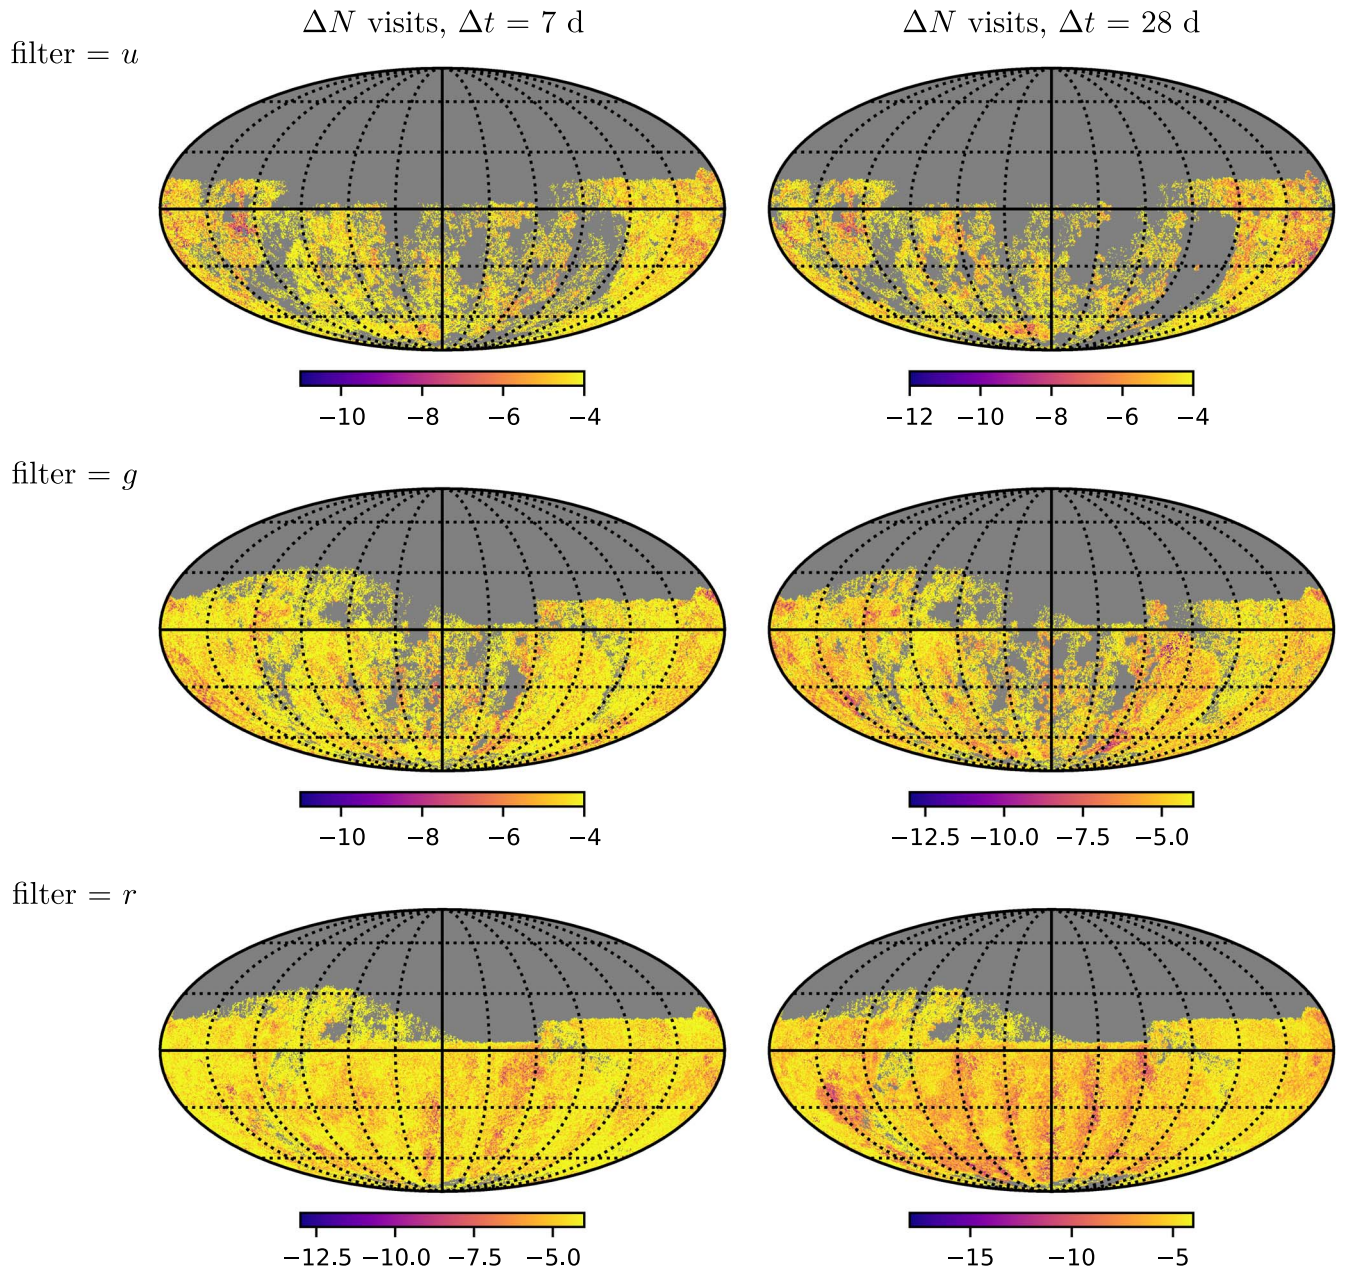


Figure 9. Sky maps ($n_{\text{side}} = 256$) showing the difference between total number of visits with templates at the end of Year 1 and the nominal `one_snap_v4.0` survey, for the u , g , and r filters. Left: template generation timescale $\Delta t = 7$ days. Right: Template generation timescale $\Delta t = 28$ days. Across most of the sky, templates are generated quickly, so the typical loss is only the four visits used to produce the templates. Compared to $\Delta t = 28$ days, the $\Delta t = 7$ day simulation has more uniform template coverage and fewer healpixels with a >4 image loss.

(The complete figure set (6 images) is available in the [online article](#).)

detected simulated objects in a given absolute magnitude bin that meets a specific metric criteria, which depend on the number and cadence of observations, coverage in certain filters, etc. We can then examine the change in these calculated metrics relative to another `rubin_sim` cadence simulation and use that output as a proxy to understand the impact of observing strategy changes on the discoverability of the simulated small body population.

To assess variations of the LSST observing strategy over the full 10 yr survey, M. E. Schwamb et al. (2023) primarily used three sets of metrics (discovery, light curve, and color light curve) that represent the top scientific priorities of the LSST Solar System Science Collaboration, reflected in their Science

Roadmap (M. E. Schwamb et al. 2018a), namely, orbit characterization, surface color/composition study, and rotational light-curve analysis. When focusing on Year 1 and the impact of different template building strategies, it is not necessarily informative to apply the exact same metrics used to evaluate the entire LSST performance when the number of visits is one-tenth of the full survey. Across most of the LSST footprint, the light-curve and color light-curve metrics applied in M. E. Schwamb et al. (2023) require more observations than will be available in Year 1, but the criteria to discover new SSOs is the same in Year 1 as it is expected to be in Year 10 of the LSST. Therefore, we apply the same discovery MAF metric applied in M. E. Schwamb et al. (2023), which examines which

Table 4

The Fractional Coverage of Unique Sky Area during Template Generation, Relative to the `one_snap_v4.0` Footprint at the End of Year 1

Δt (days)	All	<i>u</i>	<i>g</i>	<i>r</i>	<i>i</i>	<i>z</i>	<i>y</i>
3	0.960	0.593	0.796	0.921	0.891	0.848	0.910
7	0.958	0.558	0.773	0.917	0.888	0.846	0.904
14	0.955	0.545	0.751	0.911	0.881	0.840	0.900
28	0.945	0.492	0.717	0.902	0.866	0.829	0.884

Note. Results are shown per filter, and for the different template generation timescales Δt investigated in this study. This data are visualized in Figures 6 and 7 for $\Delta t = 7, 28$ days, respectively.

(This table is available in machine-readable form in the [online article](#).)

of the simulated objects meet the SSP detection criteria, i.e., three nightly pairs of observations within 15 days. In this work, we do not focus on the light-curve and color light-curve metrics as 1 yr of survey data is generally insufficient to determine these properties.

We chose to calculate the Year 1 SSO metrics for the main solar system populations: the PHAs, near-Earth objects (NEOs), MBAs, trans-Neptunian objects (TNOs), and Oort Cloud comets (OCCs) used and described in M. E. Schwamb et al. (2023). We note that the OCCs are divided into two subgroups: distant comets with perihelion (q) ≤ 20 au (labeled as OCC_r20) and those that come in closer to the Sun with $q \leq 5$ au (labeled as OCC_r5). Given the small numbers of Jupiter Trojans and 'Ayló'chaxnim types (inner Venus objects; B. T. Bolin et al. 2022, 2025) expected to be discovered in Year 1, we do not include them in our analysis. For the discovery metrics, we calculate two values for each population: one that samples the larger (brighter) objects in the population, and one for sampling the smaller (fainter) objects in the pipeline, tailoring the H thresholds to the value most representative for the given population. The full list of MAF metrics used in this work is presented in Table 3.

3. Results

For each of our template generation timescales ($\Delta t = 3, 7, 14,$ and 28 days), we analyze both the survey template coverage at the healpixel scale and the redacted Year 1 visit database described in Section 2.7, calculating relevant statistics and MAF metrics. As an overview, one can consider Figures 6 and 7, which show sky map distributions of visits in each filter for $\Delta t = 7$ and 28 days, respectively.

3.1. Template Quality and Number of Observations Used in Template Production

In Figure 8 and its accompanying online figure set, we plot, per filter, a histogram of the number of observations used to build each template and the resulting distribution of mean limiting magnitudes and the mean observed seeing values that were used to produce each of the incremental templates per healpixel. At the healpixel level, our template generation algorithm only produces a template each generation time step if there are at least four or more images with consistent quality, and a template had not been previously generated on an earlier day. We find that, across all time steps we tested, the majority of templates are created with four or five observations. This is highlighted further in Figure 9, which shows the difference in the number of visits between `one_snap_v4.0` and the

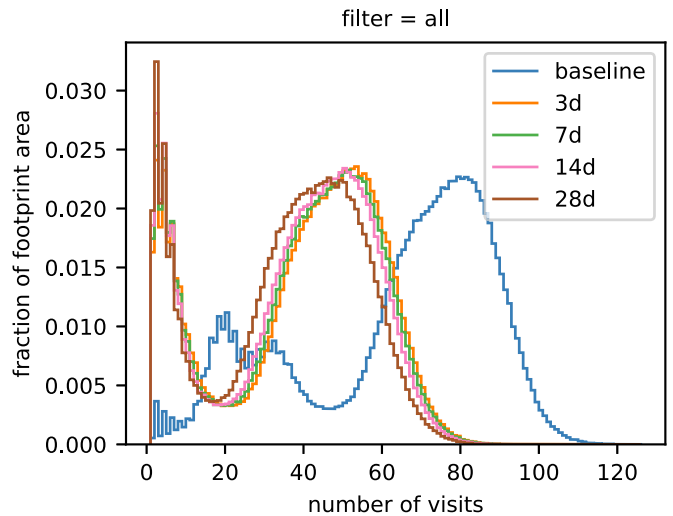


Figure 10. Histograms showing the fractional sky area (y-axis) for which healpixels have a given number of visits (x-axis) at the end of Year 1 considering all filters. We show the histogram distribution for the baseline where template generation is assumed (blue). The distributions when template generation is required ($\Delta t = 3, 7, 14, 28$ days) are denoted by the colored lines.

$\Delta t = 7, 28$ day simulations as a series of sky maps (Figures 2, 6, and 7 respectively). Most of these “missing” visits were used to create template images (Section 3.2). At longer template generation timescales, Δt , an increasing number of templates are created from >4 observations given that more images are available due to longer gaps between template building dates. Overall, the distributions of mean seeing and mean image depth are consistent across all Δt , and are comparable to the expected image quality statistics in the Year 1 baseline survey (see Tables 1 and 2). These results are not necessarily surprising given the Rubin scheduler is designed to choose a visit’s filter based in part on the observing conditions, in order to ensure the science requirements for the image quality per filter in individual and coadded observations are met (Ž. Ivezić & The LSST Science Collaboration 2013). We have applied very broad requirements on the template image quality for this study. We expect that implementing more stringent thresholds on image quality would lead to greater differences in template quality as a function of Δt . Furthermore, we would expect template coverage to increase at a lower rate as more time is needed to collect high-quality images (as considered further in Section 3.2). This analysis shows that, on average, building templates as soon as the images are available uses images with mean seeing and depth comparable to the expected Year 1 image quality. We have assumed these template images will enable RPP and SSP to produce solar system discoveries and alerts; however, a further investigation into the quality of observations required for suitable incremental templates is needed.

3.2. On-sky Template Coverage

Our analysis shows that the requirements of template generation during Year 1 of LSST reduce the number of visits, and area of the sky, with which detections can be made via difference imaging. This is shown in Figures 6 and 7, which are a series of sky maps showing the number of visits with templates for each filter, at the end of Year 1 in the $\Delta t = 7, 28$ day simulations, respectively. In comparison to the baseline scenario shown in Figure 2, the number of visits is reduced, and

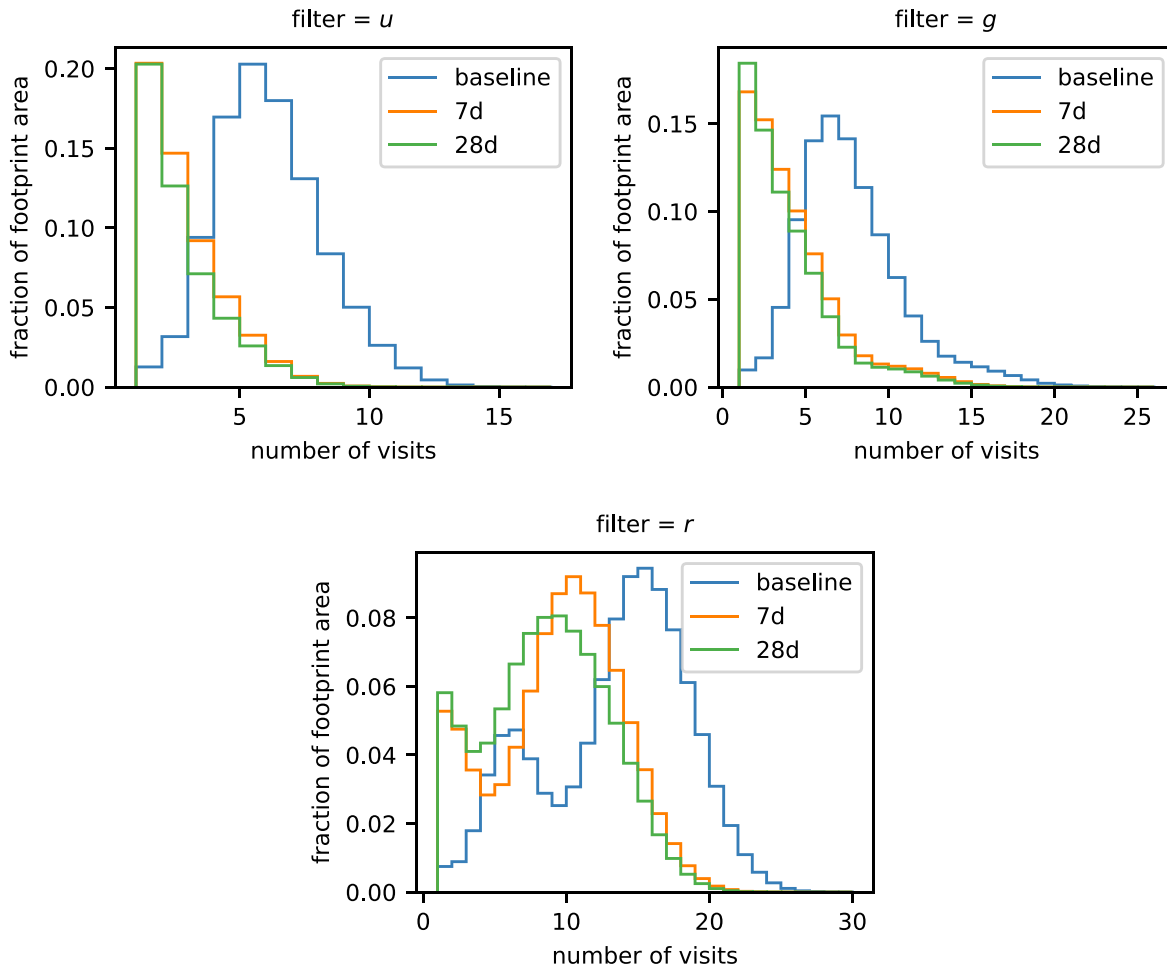


Figure 11. Per filter histograms showing the fraction of sky area with a given number of visits per healpixel (similar to Figure 10) for $\Delta t = 7$ days (orange) and $\Delta t = 28$ days (green) compared to the sky area covered in the `one_snap_v4.0` observing strategy (blue). Results are shown here for the u , g , and r filters. We compare these two timescales as there is little difference between the distributions for $\Delta t = 3, 7$ days and $\Delta t = 14, 28$ days, respectively. Each figure set element compares the baseline to a single timescale histogram for all filters.

(The complete figure set (4 images) is available in the [online article](#).)

for areas of the sky not covered by the WFD cadence, the coverage is noticeably less uniform. The fractional coverage of the survey footprint at the end of Year 1 for each filter and Δt is given in Table 4. This is summarized in Figures 9, 10, and 11, which show how the numbers of visits with templates changes relative to `one_snap_v4.0`. On average, there is a shift of $\gtrsim 4$ visits, which were required to make the templates, but this is slightly larger for the u and g filters in particular. Altogether, these results show that the loss of visits due to template generation depends strongly on the filter and area of the sky being considered.

We further investigate the dependence of template generation on region of the sky and the filter by examining selected areas of the sky in more detail. In Figure 12, we show sky maps for the `one_snap_v4.0` and $\Delta t = 7$ day simulations, and the difference between the two, in both the g and r filters. We highlight the locations representative of the nominal WFD cadence, and also the NES, which is scheduled to receive a smaller number of visits. Cutouts of these regions are shown in Figures 13 and 14 for the g and r filters, respectively. It is clear that, for both filters, the coverage in the NES is less uniform compared to the WFD when template generation is required. This effect is more extreme for the g filter compared to the r filter.

We also consider the number of visits with templates as a function of the survey time. This is shown in Figure 15, where we have plotted the cumulative sky area with templates as determined on a healpixel basis using the MAF metric (`reduceCount`) described in Section 2. We could not consider the number of visits with templates directly; each visit contains multiple healpixels (~ 183 for $n_{\text{side}} = 256$), and not all healpixels within a visit necessarily have templates due to the overlap of previous dithered visits (Figures 5 and 16). In Figure 15, the cumulative healpixel area is the nonunique area with templates, i.e., we have summed the area with templates for each visit. Therefore, this parameter is closely related to the number of visits with templates (if one were to account for the changing fractional template area per visit, Figure 16). We have considered each filter separately and included the cumulative healpixel area that would have been covered in the baseline simulation if templates existed for the full sky (black curve in Figure 15). All filters lag behind the baseline with a time delay of approximately 50 days before the template coverage begins to increase significantly, and alerts start to flow. Due to this lag, and the loss of observations required for template building, the *rizy* filters only reach 64%–69% of the baseline area for a 3 day template generation

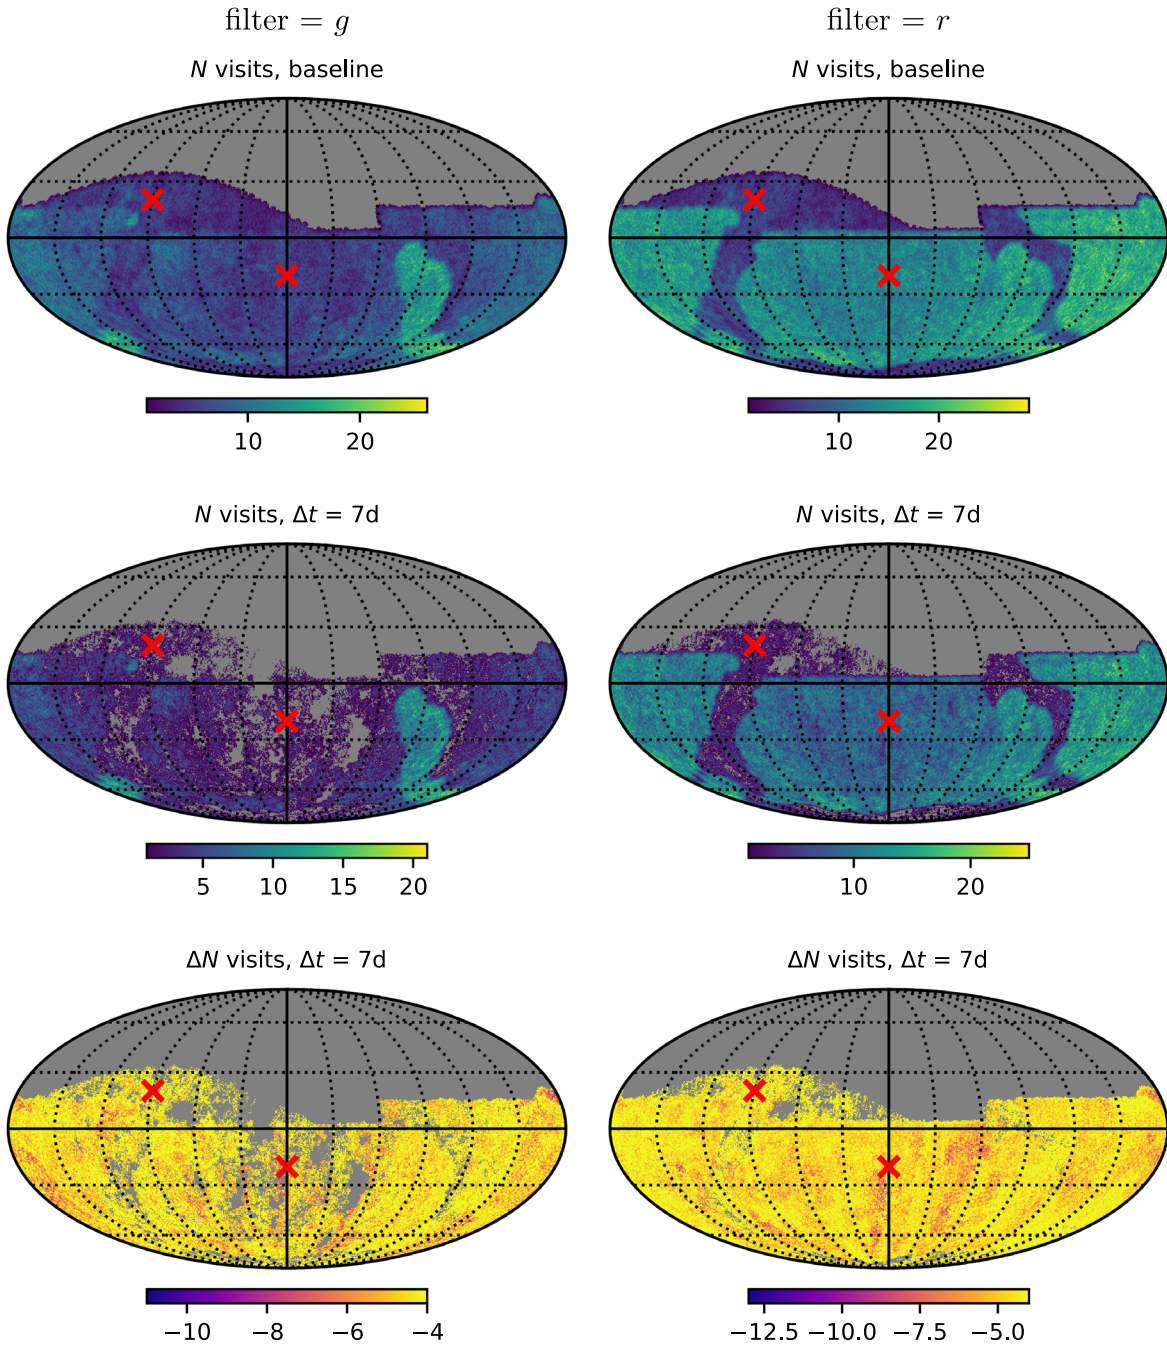


Figure 12. Sky maps for template generation timescale $\Delta t = 7$ days showing the number of visits at the end of Year 1 in different regions of the sky for the g (on the left) and r (on the right) filters. Sample positions representative of the NES and WFD (R.A., decl. positions of $(90^\circ, 20^\circ)$ and $(0^\circ, -20^\circ)$, respectively) are indicated with cross markers. The top shows the sky map for the `one_snap_v4.0`, middle shows the number of visits with templates for $\Delta t = 7$ days, and bottom displays the difference between the baseline and $\Delta t = 7$ day sky maps (all considering only the r filter). Note the different color scales for each panel indicating either the number of visits or difference in number of visits. Zoom-in plots of the regions marked with an “x” are shown in Figures 13 and 14. (The complete figure set (4 images) is available in the [online article](#).)

timescale. The u and g filters achieve noticeably lower cumulative healpixel areas than the other filters (26% and 42%, respectively, for $\Delta t = 3$ days). We note that as the template generation timescale is increased the cumulative healpixel area also decreases, with some filters experiencing greater decreases than others. The cumulative healpixel area for filters $uizy$ decreases by 4%–6% as $\Delta t = 3 \rightarrow 28$ days, whereas the drop for filters g and r is larger at 7%–9%.

These trends are also shown in Figure 17, where we again plot the cumulative healpixel area with templates, but now

normalized by the total Year 1 baseline area in that filter. Here, each panel shows the results for the different template generation timescales, and we can compare each filter directly; in all cases, the u and g filters lag behind the others. This implies that the observing strategy for the u and g filters in `one_snap_v4.0` is more sensitive to the delays incurred when template generation is accounted for.

Altogether, this analysis shows that the filters and sky areas that are scheduled to receive lower number of visits in the `one_snap_v4.0` planned observing strategy are the most




Figure 13. The panels show a zoomed-in cutout view (gnomic projection) of the sky map at the indicated locations in Figure 12 for the g filter and template generation timescale $\Delta t = 7$ days. Each cutout is a size of $11.4^\circ \times 11.4^\circ$, i.e., approximately 50×50 patches. For scale, we indicate the extent of the LSSTCam footprint, which has a 3.5° field of view (i.e., ~ 15 patches wide). The top panels show the number of visits for the `one_snap_v4.0` baseline, and bottom panels are the number of visits with templates for $\Delta t = 7$ days. Note the different color scales for each panel indicating either the number of visits or difference in number of visits. The complete figure set shows four images each for the g and r filters. Zoomed-in cutout view of g filter sky map.

(The complete figure set (4 images) is available in the [online article](#).)

negatively affected by template generation across all time-scales. This is particularly true for the u and g filters and the NES region of the sky. The removal of ≥ 4 visits for template generation in these filters/regions is a greater fractional loss compared to parts of `one_snap_v4.0` with larger numbers of visits scheduled in Year 1, such as the WFD areas. Overall, the decrease in the number of visits and survey coverage when template generation is considered (Figures 15 and 17) implies that there will be a corresponding reduction in the number of transient detections in the prompt data products (including alerts of high priority targets, e.g., solar system PHAs and ISOs) during Year 1 of LSST. We investigate this further in Section 3.3.

The need to accumulate sufficient past visits to generate templates will naturally introduce a time delay to when difference image detections start to appear in the LSST alert stream. This time delay is captured by the `MAF_deltaNight` metric described in Section 2.7, as shown in Figure 18. This




Figure 14. Same as Figure 13 for the r filter. Zoomed-in cutout view of r filter sky map.

(The complete figure set (4 images) is available in the [online article](#).)

metric measures the number of nights between the first visit (in a given filter) to a particular healpixel and the night on which its template (in that filter) was generated. Rather than the previous analysis of considering on what night of the survey a given healpixel had a template, this metric describes how long it takes to build a template once that particular healpixel is observable by the survey. This therefore removes the bias of which area of the sky is the current focus of the survey. Figure 18 shows that, as Δt is decreased, the time delay to create a template is also decreased. In this figure, the cumulative fraction of footprint area refers to the unique on-sky area planned for Year 1 of the survey in that filter; therefore, the figure indicates the rate at which the template coverage increases. The `deltaNight` metric in Figure 18 is more closely modulated by the choice of Δt than the `reduceCount` metric in Figure 15; this is particularly clear for the u filter results. Across all filters, the differences between different Δt are most pronounced for low values of `deltaNight` when it is less likely for there to be sufficient revisits to a given field to generate templates. As `deltaNight` $\rightarrow 90$ days, all filters have started to converge, and we see that up to $\sim 90\%$ of each filter's survey footprint has templates, as 90 days is several times greater than any Δt considered here.

As mentioned in Section 2.7, the on-sky extent of each visit will lie across multiple healpixels over which the `MAF` metrics are being calculated. Not all healpixels within a visit footprint will necessarily have a template at the time of exposure, due to

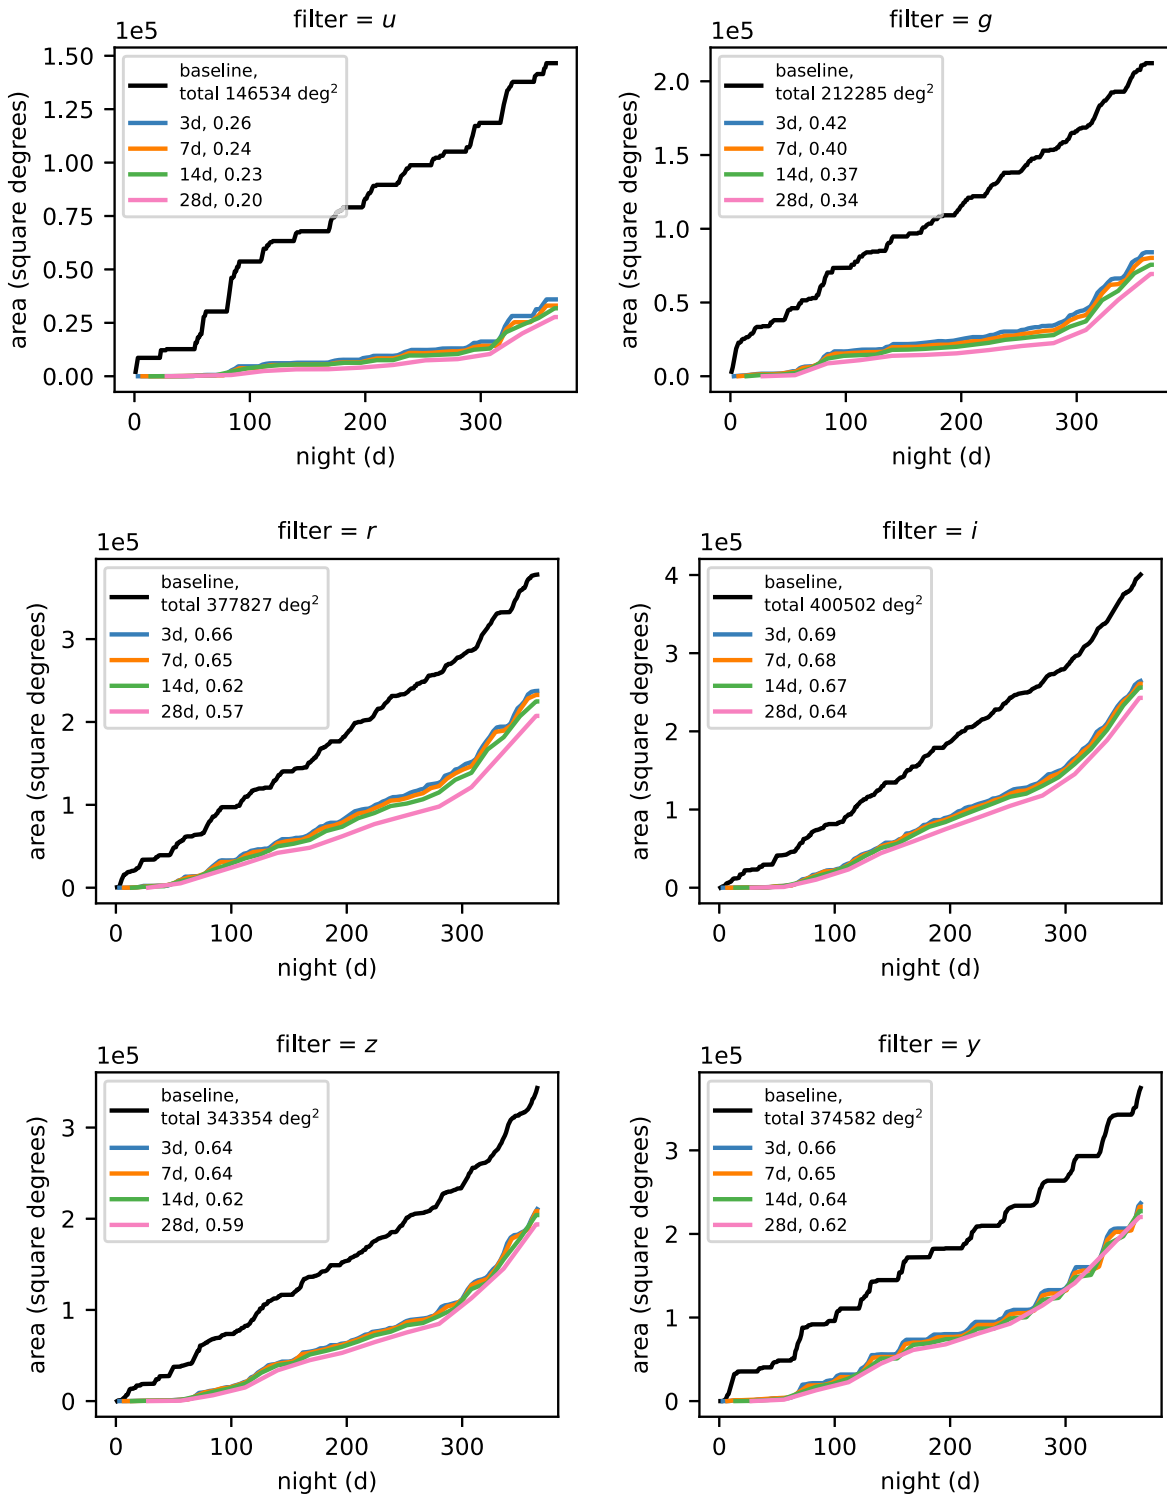


Figure 15. The cumulative healpixel area with templates for a given filter in Year 1 of LSST. Each panel shows the results for the *ugrizy* filters, respectively. The area covered by the nominal `one_snap_v4.0` survey (assuming all templates exist at all times) is shown by the black line. The colored lines show the increase in area when template generation is considered ($\Delta t = 3, 7, 14, 28$ days). The legend states the total area covered in that filter for the Year 1 baseline survey, and the fraction of that area reached for the different template timescales by the end of Year 1. The cumulative area is approximately related to the cumulative number of visits with templates given that each visit has an area of 9.6 deg^2 , but that visit may not have 100% template coverage.

the dithering/rotation of past exposures. Therefore, as described in Section 2.7, we recorded the number of constituent healpixels with templates for each visit. Figure 16 shows a histogram of the fractional template coverage for all visits in Year 1 of `one_snap_v4.0`. The large spike at zero indicates the visits that had absolutely no template coverage and so could

not have generated any alerts. The majority of these visits were subsequently used to make templates; we estimate the total number of Year 1 visits used for template generation in Table 5. Figure 16 demonstrates that the number of visits with a template coverage $\geq 90\%$ increases with more frequent template generation. For $\Delta t = 28$ days, 38.8% of Year 1

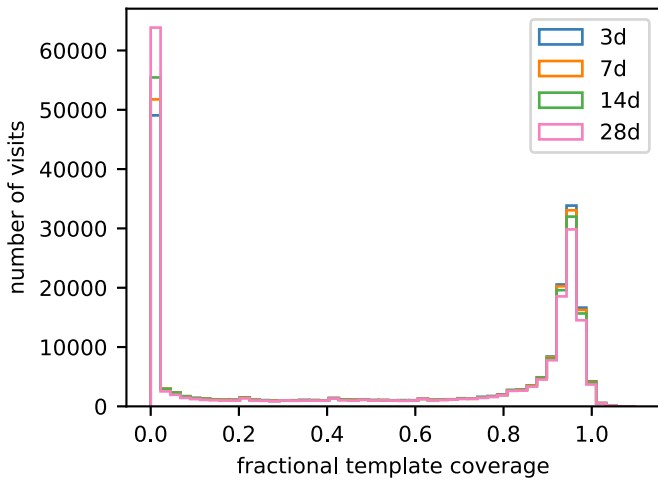


Figure 16. Histogram of the fractional template coverage for all Year 1 visits, which is determined from the number of healpixels within the visit footprint with templates. Results are shown for the `one_snap_v4.0` cadence simulation, assuming a range of template generation timescales. The peak at zero consists mainly of the images used to make the templates (see Table 5).

visits have a template coverage $\geq 90\%$; this increases to 43.6% for $\Delta t = 3$ days. Note that the small number of visits with coverage $> 100\%$ arises due to healpixels that include only part of the visit footprint; the footprint does not align perfectly with the healpixel grid, and we may overestimate area due to healpixel resolution. To justify our selection of a 90% template coverage threshold, we present how the distribution shown in Figure 16 changes as a function of time (Figure 19). This figure shows that throughout most of Year 1 there is an approximate bimodal distribution between visits with 0% and 90% template coverage. Therefore, if we were to reduce the 90% threshold in our analysis, we would not include significantly more visits, and such low-coverage visits would be less useful for a moving object discovery (see Section 3.3 for further discussion).

We also include animated figures showing the cumulative sky map visit coverage in Year 1 for the `one_snap_v4.0` simulation and the various template generation simulations. Figures 3 and 20 show snapshots of these animations for the r filter, where the template simulation is $\Delta t = 7$ days. These animations highlight the time lag when template generation must be considered, and the patchier coverage when visits without templates are rejected.

3.3. Solar System Discovery Metrics

For those observations that have sufficient template coverage at the time of the exposure, we can analyze the Year 1 SSO metrics, considering only visits that had a fractional template coverage of $\geq 90\%$, as discussed in Section 3.2 above. We note that up to this point our analysis has focused on healpixels with a size comparable to LSST patches ($13.7'$), and now, we are considering visits of the size $3.5'$. We have made the conservative assumption that the SSP pipeline will only be effective for visits with a sufficient level of template coverage. Cutting visits with $< 90\%$ template coverage will naturally remove many healpixels that did have templates; the sky maps shown previously (e.g., Figure 6, etc.) would be significantly sparser if replotted for only visits with $\geq 90\%$ template coverage rather than individual healpixels with templates.

Figure 21 shows the results for the discovery metric of objects with three pairs of detections over a 15 night period. As

described in Section 2.7, we have considered a range of dynamical populations: MBAs, NEOs, PHAs, TNOs, and OCCs, with a maximum perihelion distance of 5 au (`OCC_r5`) and 20 au (`OCC_r20`). The different orbits and physical properties of these distinct populations have a strong influence on their discovery. The metrics are divided into two components looking at the absolute magnitude bins that represent the “bright” and “faint” objects of each dynamical population (see Table 3) in order to assess the effects for large and small SSOs separately.

We present the metric results of each template generation timescale relative to the default `one_snap_v4.0` cadence simulation where the presence of templates is implicitly assumed (i.e., all visits are capable of making SSO detections and alerts). Figure 21 shows that the requirements of template generation (across all timescales tested here) will strongly impact discovery in Year 1, with decreases of tens of percent in the MAF discovery metric. The discovery of faint objects is more strongly affected; fainter objects are observed with lower signal-to-noise ratio (SNR), so there are fewer possible detections in the baseline. There is a larger fractional loss of detections due to a lack of templates compared to the bright population. Furthermore, there are large differences in discovery between different SSO populations such as the OCCs and MBAs; this is primarily due to how these objects move across the sky. The inner solar system populations generally move faster; therefore, an MBA is more likely to pass through sections of the sky without templates and not get enough detections to be discovered. In comparison, a slower moving TNO in the outer solar system covers less sky during discovery; its detection requires a smaller area of the sky having templates and is therefore more likely. For each population, as Δt increases, there are only modest decreases of a few percent in discovery, which is small compared to the overall effect of template generation.

We performed an additional analysis assessing how differences in the survey cadence covering the ecliptic plane, where most SSOs are located, affect discovery. We split the `one_snap_v4.0` survey into visits with $\text{decl.} \geq 0^\circ$, where the NES is located, and $\text{decl.} < 0^\circ$, where the southern part of the ecliptic is primarily sampled by the WFD cadence. Following the previous methodology, these sets of visits were analyzed separately by the MAF SSO discovery metrics, the results of which are shown in Figure 22. For visits with $\text{decl.} < 0^\circ$, there is an increase in the fraction of SSOs discovered during template generation compared to the previous analysis of the whole sky. Accordingly, the drop in discoveries is far greater for the $\text{decl.} \geq 0^\circ$ visits. This demonstrates that losses in the NES region are dominating the overall reduction in SSO discoveries during Year 1, whereas the areas of the ecliptic sampled by the WFD are not as severely impacted. In Year 1, the NES receives $\sim 75\%$ fewer visits than the WFD; this means that losing ≥ 4 visits to templates is a greater proportional loss, which results in a larger fraction of missed SSO detections. Furthermore, we have only considered SSO discovery for visits with $\geq 90\%$ template coverage rather than for all healpixels with templates. This means that the level of the coverage in the NES shown in Figures 6, etc. is actually much lower when only visits with $\geq 90\%$ template coverage are considered. Only a small number of visits in the NES meet the template requirements, which explains the strong decrease in the discovery metric for this region.

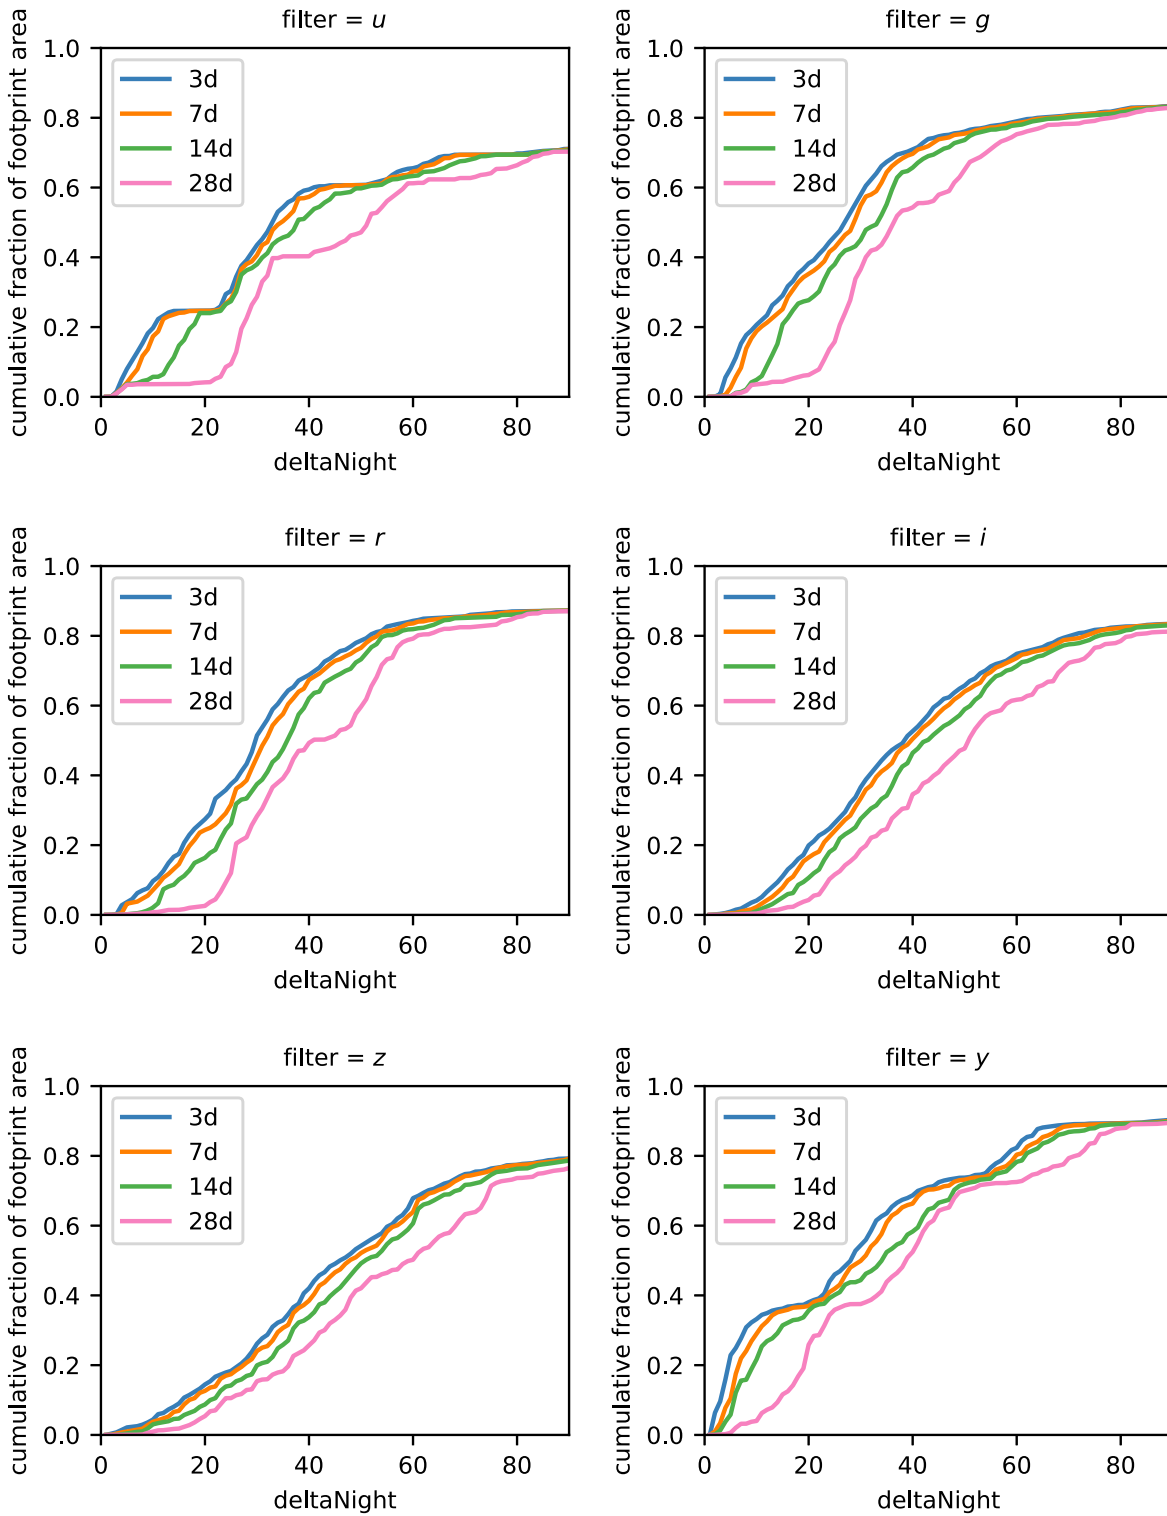


Figure 17. An alternative to Figure 15, where now we plot the cumulative area covered in a particular filter as a fraction of the total baseline area in that filter at the end of Year 1. In each panel, we show the results for $\Delta t = 3, 7, 14, 28$ days, and the color of each curve indicates the filter (*ugrizy*). The legend provides the fractional cumulative area compared to the baseline at the end of Year 1 for each filter.

In Figure 23, we consider in more detail how the cumulative completeness (that is the fraction of objects in the sample that have been discovered) increases as a function of the survey time. Compared to the baseline, discoveries for all populations (and all Δt) are delayed by approximately 70 days. In other words, our simulation of template generation implies that there

will be no SSO discoveries in the first couple months of the survey. This delay arises primarily from the ~ 50 day timescale required to build up sufficient visits for a given patch of the sky (Figure 15), plus the additional time to acquire the detections for SSO discovery (three nightly pairs over a ~ 2 week period). Because of this temporal offset relative to the baseline survey,

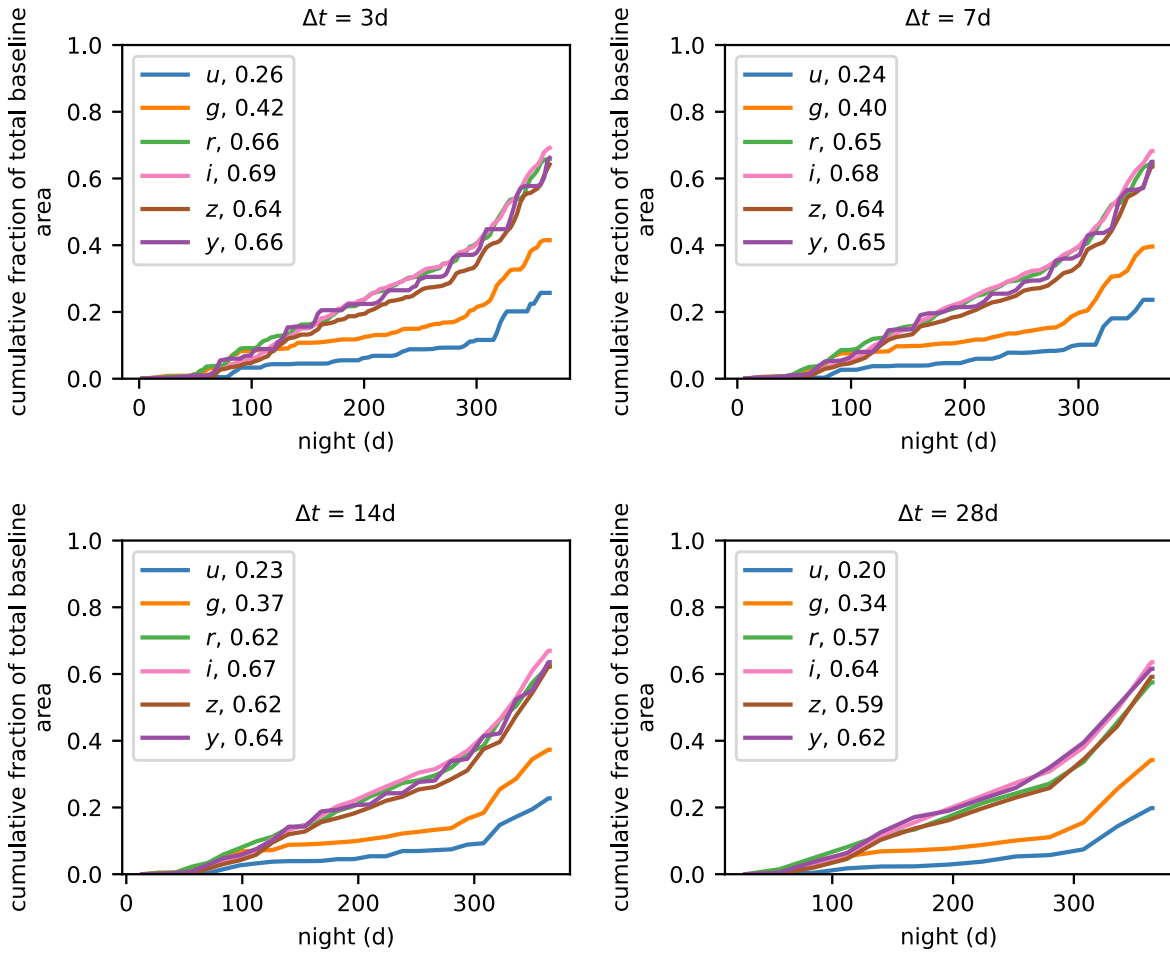


Figure 18. Plots showing the cumulative fraction of the Year 1 survey footprint area (i.e., the unique on-sky area covered by that filter in the Year 1 survey) against `deltaNight`, which is the number of days between the first visit to a healpixel and the night at which a template was generated. Each panel shows the results for a different filter (*ugrizy*). The *x*-axis is limited to show a maximum `deltaNight` = 90 days to highlight the timescale over which there is the greatest change.

Table 5
The Number of Visits that Were Used to Generate Templates for Each Filter for Different Template Generation Timescales Δt

Δt (days)	All	<i>u</i>	<i>g</i>	<i>r</i>	<i>i</i>	<i>z</i>	<i>y</i>
3	68,572	9596	11,606	12,234	11,709	11,398	12,030
7	71,355	10,034	12,078	12,758	12,053	11,676	12,755
14	74,584	10,163	12,624	13,726	12,657	12,126	13,288
28	80,656	10,732	13,273	15,347	14,007	13,181	14,115

Note. As we conducted the template analysis on a healpixel level, this is an approximate number of visits, calculated as number of visits used to generate the first template in all healpixels times healpixel area ($5.25 \times 10^{-2} \text{ deg}^2$), divided by the camera footprint area (9.6 deg^2).

(This table is available in machine-readable form in the [online article](#).)

the total number of discoveries at the end of Year 1 is reduced across all test populations, similar to the results in Figure 21.

Accounting for this delay, the cumulative discovery curves approximately track the baseline; however, there are times when discoveries plateau in our simulations. This is because our analysis depends strongly on the survey start date and when the ecliptic, where a significant number of SSO are located, is visible. Based on previous versions of the Rubin Project timeline, the start date for `one_snap_v4.0` was assumed to be 2025 May. From the observing site, the NES is observable by the survey for the first time in the period 2025 September–February. Further delays to the survey start date could reduce

the amount of time available to observe the NES and have an even more dramatic effect on SSO discoveries. Future investigations should account for any changes to the survey schedule as the timeline becomes more clear.

The results of these MAF solar system metrics provide a more direct estimate of the effects of template generation on solar system science than simply looking at the sky area with templates considered in Section 3.2. We caution the reader that these MAF metrics are for the total numbers of discovered SSOs, per simulated bin in absolute magnitude. It would require additional assumptions to convert the metrics into an absolute total number of objects, so we instead continue to

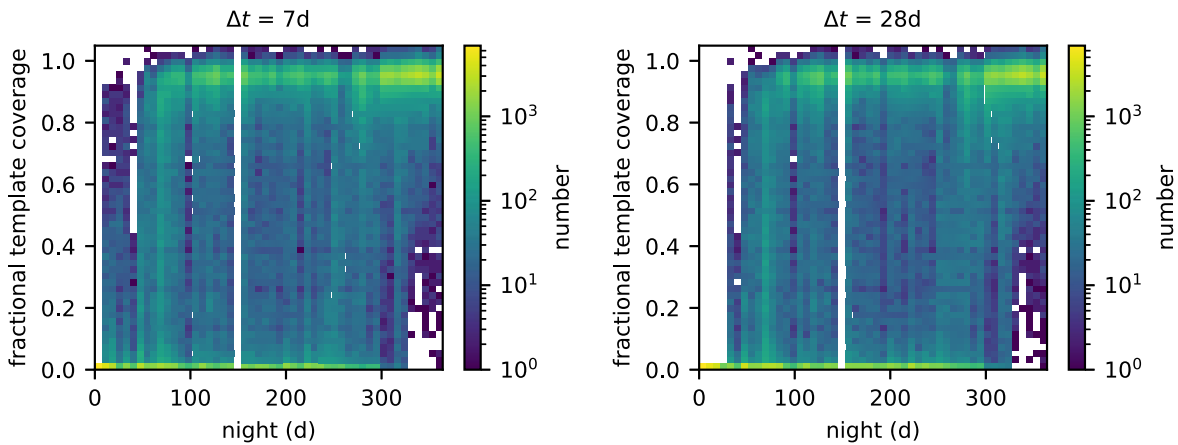


Figure 19. Two-dimensional histogram distributions showing how the fractional template coverage of all visits changes as a function of survey time. The logarithmic color scale indicates the number of visits with a particular template coverage on a given date. Results are shown for $\Delta t = 7$ days (left) and $\Delta t = 28$ days (right). Empty spaces indicate survey downtime when no visits were taken, e.g., scheduled maintenance around night 150.

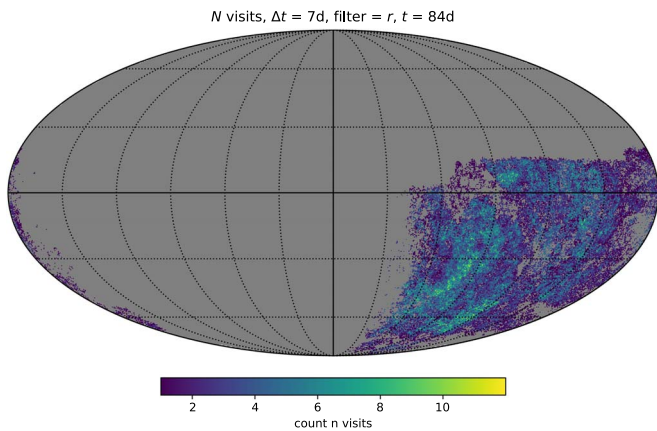


Figure 20. A snapshot of a video animation of the sky coverage over time in `one_snap_v4.0` observing strategy in the case where templates are produced incrementally every $\Delta t = 7$ days over the first year of the simulated survey. The animation steps through the first year of the simulated LSST in intervals of 7 days, displaying the cumulative number of on-sky visits in the r filter. The plots are centered on R.A. = 0° and decl. = 0° . R.A. and decl. lines are marked every 30° .

(An animation of this figure is available in the [online article](#).)

examine the relative change to the baseline. We can clearly see that the requirements of template generation lead to large drops in discovery for the dynamical populations investigated. If we consider the analysis in M. E. Schwamb et al. (2023), they compared different survey strategies and stated that a $\pm 5\%$ variation in a given metric (relative to the baseline survey at the time) would be acceptable for solar system science cases. However, these results show drops of 28%–63% in discovery relative to the baseline across all test populations (Figure 21), alongside delays of $\gtrsim 60$ days before SSO discoveries can be made (Figure 23). Altogether, these results highlight that the number and rate of “real-time” solar system discoveries in Year 1 will be significantly impacted by an incremental template generation strategy that builds templates from the regular LSST observations as the survey progresses. The NES is expected to be the hardest hit region within the LSST footprint. This will affect the discovery of SSOs that require a rapid follow-up, such as ISOs, PHAs, and minimoons.

4. Summary and Conclusions

We have investigated the impact of incremental template generation in Year 1 of LSST using the `rubin_sim` simulated `one_snap_v4.0` observing strategy and associated LSST MAF metrics. Focusing on solar system discovery metrics, we used MAF to simulate the effects of generating templates in the first year of the survey. This was done by dividing Year 1 into a series of template generation nights and counting the number of visits in each healpixel of the sky and establishing when the criteria for template generation was met. We utilized only observations from the main survey, i.e., excluding DDF and low-solar elongation twilight visits, assuming that there was no significant contribution of template-grade observations made during the commissioning of Rubin Observatory. Having used a $\Delta t = 3, 7, 14,$ and 28 day template building cadence coupled with the requirement of $\geq 90\%$ template coverage for each LSST observation, we summarize our findings below:

1. Early template generation is preferred. The sooner templates are generated, the closer the performance will be to the idealized `one_snap_v4.0` observing strategy’s real-time discovery rates. This is because at shorter template production timescales more observations become available for nightly prompt processing. We expect it will take ~ 50 days from the start of the survey to build significant template coverage across the sky, followed by ~ 2 weeks before SSO discoveries are made by the SSP pipeline. This results in a total delay of over 2 months before significant numbers of nightly SSO alerts will be issued. Across all the metrics we evaluated, a monthly template strategy performs the worst. A 3 day/7 day turnaround time for template production provides the most opportunities to identify transient sources in LSST observations in real time and results in more uniform coverage across the sky.
2. Solar system discoverability metrics are significantly affected by template generation, with the shorter template generation timescales performing slightly better. The discovery metrics across all dynamical populations tested here have a $\gtrsim 28\%$ decrease compared to the nominal Year 1 baseline metrics. The MBA discovery metrics are impacted the most with a drop $> 40\%$ for the larger

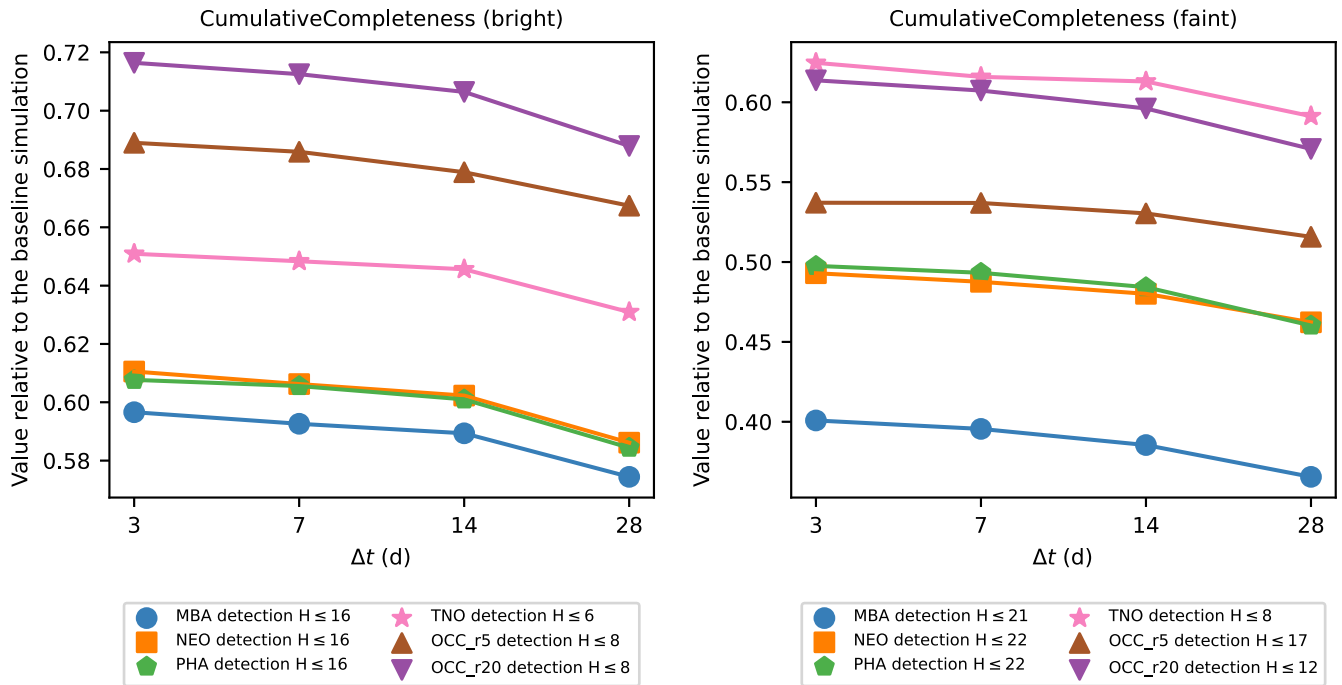


Figure 21. Here, we show the results of running the MAF SSO discovery metrics on a redacted Year 1 visit database of visits with template coverage $\geq 90\%$. This metric requires three detection pairs over the space of 15 nights. We consider several dynamical populations, which are indicated by color and marker shape. The results of this metric are shown for $\Delta t = 3, 7, 14, 28$ days and are given relative to the results for the `one_snap_v4.0` Year 1 survey in which templates are assumed to already exist. The left panel shows the discovery metric for the “bright” objects in each population (low H), and the right panel shows the “faint” objects (high H) as defined in Table 3.

($H \leq 16$) objects and $>60\%$ for the smaller asteroids ($H \leq 21$).

- For all incremental template timescales, the coverage in u -band and g -band visits lags significantly behind the other filters due to the small number of scheduled observations in these filters during Year 1. Likewise, due to the reduced number of visits planned in the NES (compared to the WFD), a significant number of Year 1 NES visits will not have templates and will not be processed by the nightly RPP pipelines for the identification of transient sources and moving SSOs. The overall loss of SSO discoveries in Year 1 is dominated by the lack of templates in the NES. The discovery metrics within the NES are half the value computed for WFD fields.

Without modifications to the baseline (`one_snap_v4.0`) Year 1 observing strategy, there will still be some LSST observations with suitable templates made as the survey progresses to allow image subtraction and alert generation within 60 s of the image readout, enabling the RPP and SSP pipelines to run and generate nightly solar system discoveries/alerts. But there is a nonnegligible hit to the real-time detection of small solar system bodies. The change in MAF metrics is much higher than the 5% difference from the baseline strategy deemed acceptable in the evaluation of cadence parameters for optimizing LSST solar system science by M. E. Schwamb et al. (2023). For each simulated small body population investigated here, the larger objects are impacted less because they are typically brighter in apparent magnitude and thus more likely to be detectable in most or all of the LSST observations that they are present in. Smaller (and typically fainter) objects require more chances to be able to detect them, and thus are impacted more severely by incremental template generation strategies.

The drop in the discovery metrics is likely influenced by the lack of g -band templates, as g is typically paired with the r filter in nightly pairs. Furthermore, the lack of templates in the NES region leads to a significant loss of SSO discoveries. Exploring options for boosting Year 1 template production in low-coverage filters and sky regions (i.e., the g filter and NES) is likely the first step to improving the number of nightly pairs searchable each night for moving objects.

In this work, we have focused on the impact on solar system metrics, but there are parallels that can be drawn to the discovery of other astrophysical transients using the Rubin alert stream. An LSST visit will only be able to produce alerts when the visit’s on-sky footprint has adequate template coverage, which requires a sufficient number of previous observations (in the same filter). In this study, we have demonstrated that template generation reduces the area of survey sky able to generate alerts, focusing on the impact on SSO science. However, our results imply there will be reductions in real-time astrophysical transient alerts across the science cases in Year 1. For example, it will take several observations to identify the characteristic brightening in a transient light curve consistent with a supernovae or tidal disruption event. Additionally, the variable star and astrophysical transient community has advocated for different filters to be used in LSST nightly pairs (F. B. Bianco et al. 2019), with u and g bands paired with r to explore color evolution over a night (Ž. Ivezić & The SCOC 2021; F. Bianco & The SCOC 2022). We have demonstrated that the template coverage of bluer u and g filters lags behind the redder filters. This will impact the availability of pairs of detections, therefore reducing chances to discover and characterize fast-changing explosive transients. Furthermore, we expect that the lower number of scheduled visits in the Galactic plane will lead to a reduced number of templates

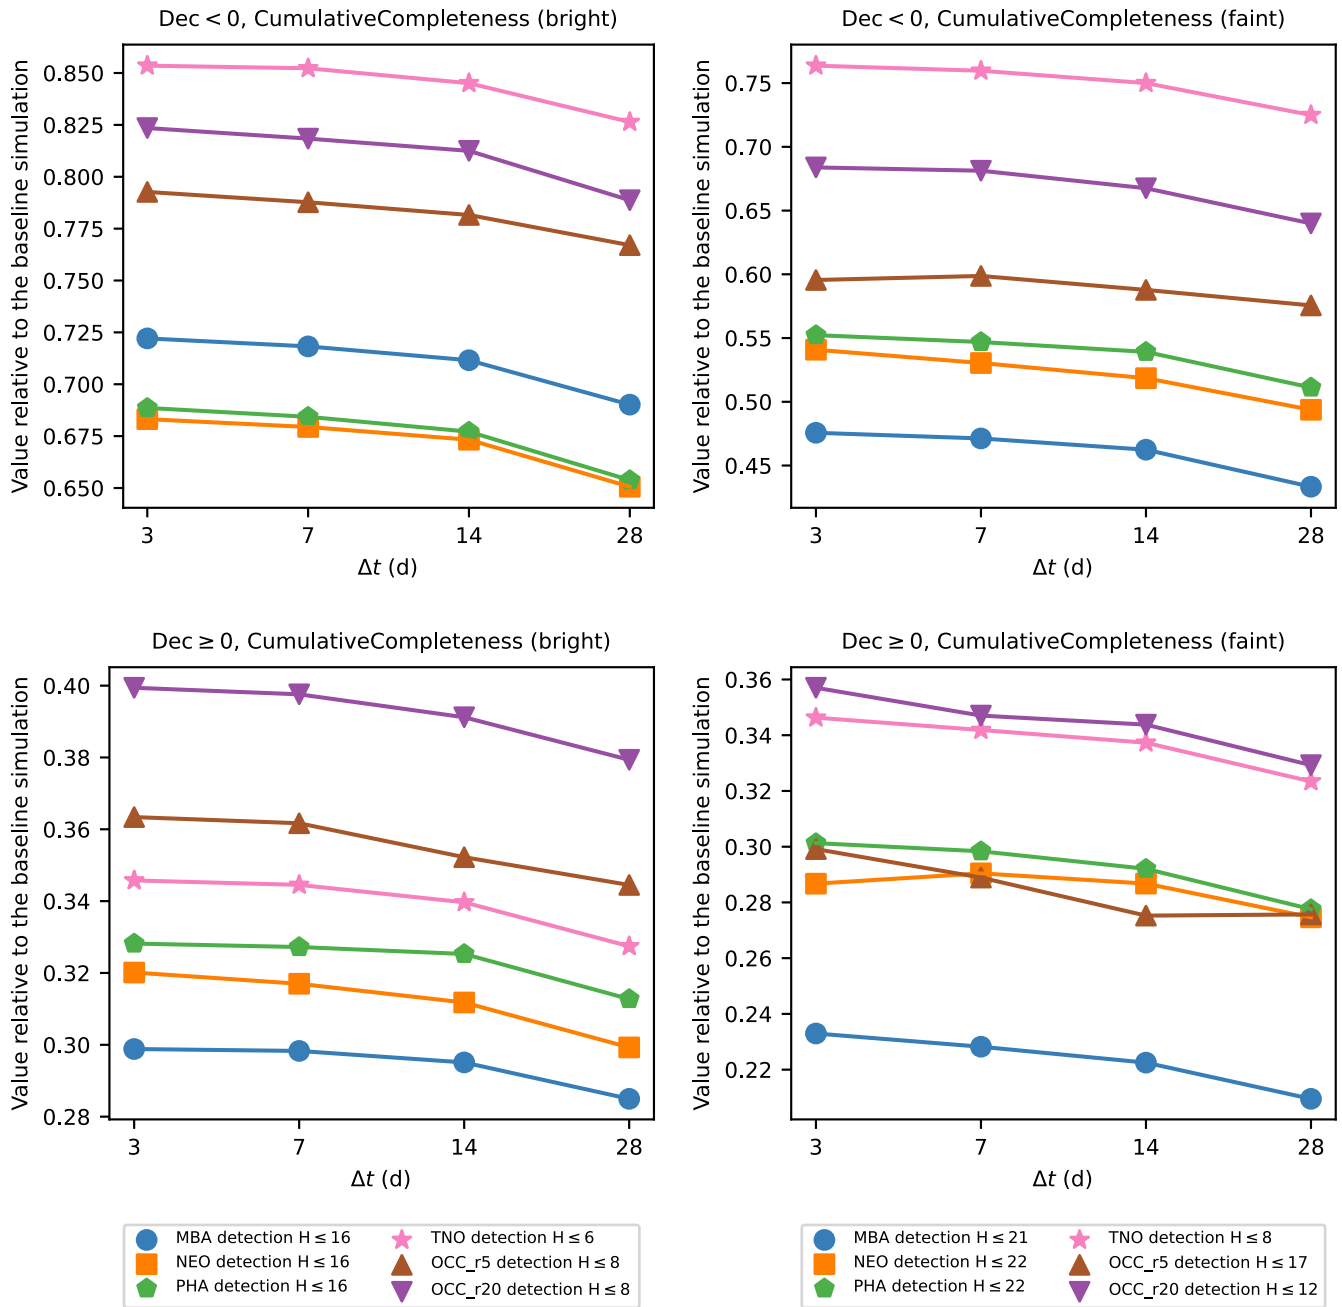


Figure 22. These panels show the same discovery metrics as Figure 21, but now, results are for subsets of visits with decl. < 0° and decl. ≥ 0° (upper and lower panels, respectively).

and detections in this region, similar to what we have observed for the NES for SSOs. Further work is needed to explore the full impact of incremental template generation on the discovery and rapid follow-up of transients in the various LSST science goals during Year 1 of the survey.

This work is a first step toward understanding the expected discovery yields and discovery output from the RPP data products and Rubin alert stream in Year 1 operations. The Rubin Observatory operations and data management teams have yet to settle on an incremental template strategy (M. L. Graham et al. 2020; L. P. Guy et al. 2023). Based on our analysis, we strongly recommend that they explore the possibility of generating templates on weekly or shorter timescales. We have assumed that there is no significant contribution to Year 1 templates from commissioning data and

that the `one_snap_v4.0` observing strategy is executed with no modifications to prioritize incremental template building during Year 1. We have also assumed that the first four suitable observations of a given patch are used to produce templates with generous image quality constraints. Compared to later templates made from more images, these Year 1 incremental templates may have more artifacts and/or lower SNR, which could impact the detection of faint sources or the identification of extended sources (which are an indicator of cometary activity for example) against the background noise. A further analysis is needed to explore more sophisticated strategies and differing requirements for incremental template generation.

The LSST observing cadence has been carefully optimized to maximize results for all main science goals across the 10 yr survey; however, it is not necessarily optimized for incremental

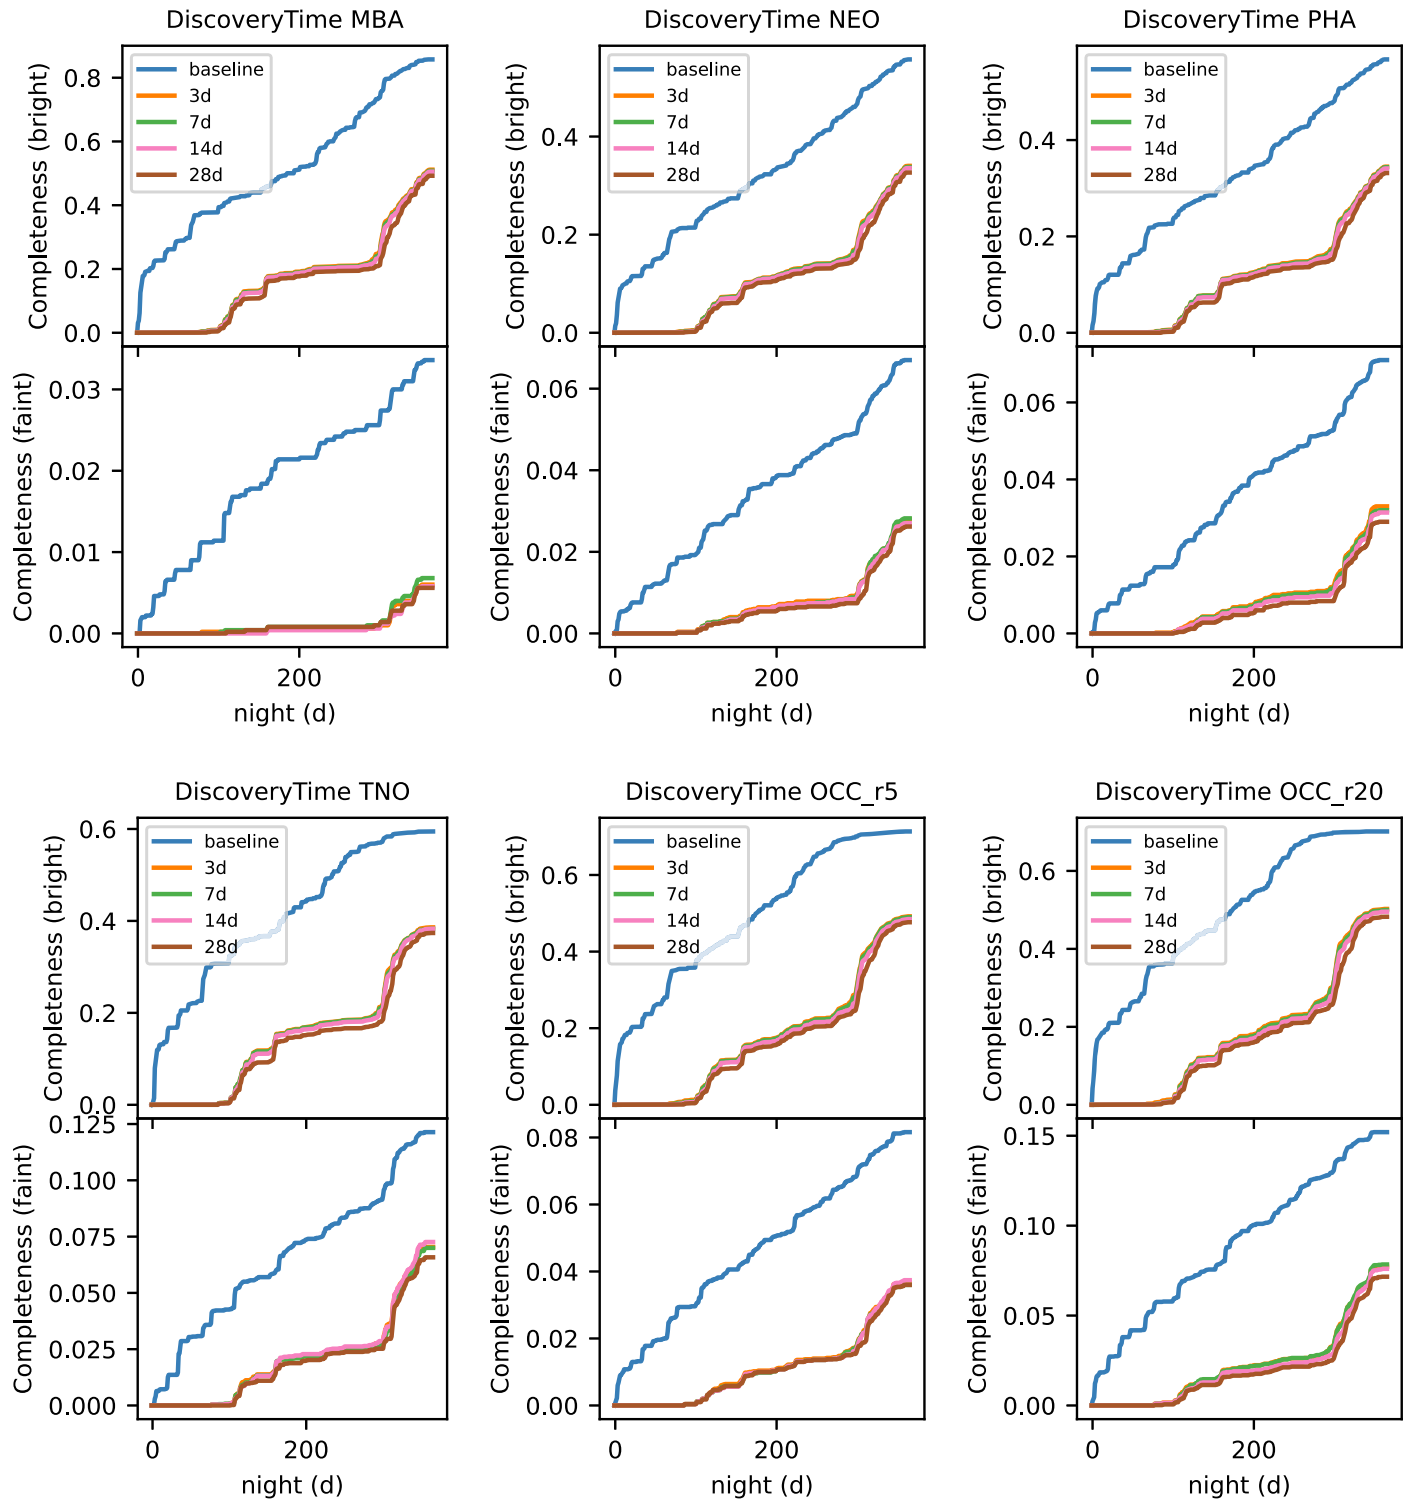


Figure 23. Plots showing the discovery completeness of each population during Year 1 of the survey, i.e., the cumulative fraction of objects that have been discovered as a function of time. The line colors distinguish the `one_snap_v4.0` survey where template generation has been implicitly assumed and the simulations in which template generation has been considered (for $\Delta t = 3, 7, 14, 28$ days). The various panels display results for the different bright/faint dynamical populations considered in this work (see Table 3).

template generation in Year 1. This is because the alerts “lost” due to template generation in Year 1 will be retrieved when all data are reprocessed at the first data release; but we lose the “real-time” nature of the transient alerts and the opportunities for follow-up in Year 1. This could result in extremely rare scientific discoveries, such as a briefly observable ISO, being

missed in Year 1. Further `rubin_sim` simulations are needed to examine if, and how, the Rubin scheduler can be tuned to maximize template production over the Year 1 without significantly impacting the 10 yr outputs of the survey. It is important for the community and the SCOC to explore possible modifications to the Year 1 observing cadence that may

significantly boost template production. For SSO science in particular, possible methods of boosting the number of observations in the NES in Year 1 should be explored (while considering any impacts on the wider survey science goals). This might be achieved by reallocating an amount of Year 1 observing time to focus on template building. For example, regions lacking template coverage (namely, the NES) could be prioritized by the scheduler when observable, by taking time from regions with more than sufficient coverage (e.g., the WFD). Likewise, an amount of time could be shifted from well-observed filters into other filters, such as u and g . Furthermore, one could consider “borrowing” NES visits from later years of the survey and pushing them into Year 1 instead. This would improve the coverage of the NES in Year 1 at the expense of slightly reduced numbers of visits in later years.

It is vital that studies exploring the possible parameter space for incremental template generation within the LSST observing strategy and its impact on the LSST 10 yr and Year 1 MAF metrics be completed before the start of Rubin on-sky commissioning with LSSTCam (currently expected to start in the first half of 2025). This will provide an opportunity to optimize the Year 1 observing strategy for alert production while giving the community time to initiate follow-up preparations for the start of LSST with the best picture of what Year 1 nightly outputs of the survey will look like.

Acknowledgments

This work was supported in part by the LSST Discovery Alliance Enabling Science grants program, the B612 Foundation, the University of Washington’s DiRAC (Data-intensive Research in Astrophysics and Cosmology) Institute, the Planetary Society, Karman+, Breakthrough Listen, and the Adler Planetarium through generous support of the LSST Solar System Readiness Sprints. The DiRAC Institute is supported through generous gifts from the Charles and Lisa Simonyi Fund for Arts and Sciences and the Washington Research Foundation. Breakthrough Listen is managed by the Breakthrough Initiatives, sponsored by the Breakthrough Prize Foundation (<http://www.breakthroughinitiatives.org>). J.E.R. acknowledges support via the Science Technology Facilities Council (STFC) funding for UK participation in LSST, through grant ST/X001334/1, in addition to support from the Royal Society RF\ERE\231044. M.E.S. and S.R.M. were supported by the UK STFC grants ST/V00691/1 and ST/X001253/1. M.E.S. also acknowledges the support by a LSST Discovery Alliance LINCC Frameworks Incubator grant [2023-1042SFF-LFI-01-Schwamb]. Support was provided by Schmidt Sciences. M.J. acknowledges the support from the University of Washington College of Arts and Sciences, Department of Astronomy, and the DiRAC Institute, the Washington Research Foundation Data Science Term Chair fund, and the University of Washington Provost’s Initiative in Data-Intensive Discovery. C.O.C., S.G., and P.Y. acknowledge support from the DiRAC Institute in the Department of Astronomy at the University of Washington. The work of S.G. is supported by NOIRLab, which is managed by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the U.S. National Science Foundation. The work of S.R.C. was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (#80NM0018D0004).

This material or work is also supported in part by the National Science Foundation through Cooperative Agreement AST-1258333 and Cooperative Support Agreement AST1836783 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.

We acknowledge access to `cuillin`, a computing cluster of the Royal Observatory, University of Edinburgh. The authors thank Eric Bellm, Leanne Guy, and Kat Volk for useful discussions. We thank Federica Bianco and the American Astronomical Society (AAS) journals editorial team for facilitating the Rubin LSST Survey Strategy Optimization *ApJS* focus issue. This research has made use of NASA’s Astrophysics Data System Bibliographic Services. Some of the results in this work have been derived using the `healpy` and `HEALPIX` packages. For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) license to any author accepted manuscript version arising from this submission.

Data access. Data used in this paper are openly available from the Vera C. Rubin Observatory Construction Project and Operations Teams. The `rubin_sim` and `rubin_scheduler` LSST cadence simulation databases used in this work are publicly available from P. Yoachim (2024). The databases of template coverage per visit derived in this work are available on the Edinburgh DataShare at doi:10.7488/ds/7893.

Facility: Rubin:Simonyi

Software: `healpy` (K. M. Górski et al. 2005; A. Zonca et al. 2019), LSST Metrics Analysis Framework (MAF; R. L. Jones et al. 2014), `Matplotlib` (J. D. Hunter 2007), `rubin_sim` (A. J. Connolly et al. 2014; F. Delgado et al. 2014; LSST Science Collaboration et al. 2017; E. Naghib et al. 2019; R. L. Jones et al. 2020), `pandas` (W. McKinney 2010), `SciPy` (P. Virtanen et al. 2020), `numpy` (C. R. Harris et al. 2020).

Author Contributions

J.E.R. developed the code to simulate template generation based on the MAF framework. He ran the simulations, made plots, and interpreted results. He provided the primary text for methods and results sections, which were edited mainly by M. E.S. He provided comments, feedback, and edits on sections written primarily by M.E.S.

M.E.S. contributed to the discussions about the implementation of the incremental templates to the LSST observing strategy cadence simulation and the choice of MAF metrics to review. She was also heavily involved in the discussions with J. E.R. about the analysis and interpretation of the results. She contributed text to primarily the first half of the manuscript and the summary and conclusions section (Section 4) that were later edited and revised by J.E.R. She also made Figure 4 and Table 3. She also gave feedback on the manuscript’s figures. She also provided feedback on the overall paper draft.

R.L.J. provided significant input into how to account for incremental templates to the LSST cadence simulation and apply MAF metrics. She also developed the initial notebook for looking at filter completion for visits that were utilized in this work.

R.L.J. and P.Y. provided valuable example code and advice from which the project was developed. R.L.J. and P.Y. also provided guidance on the Jupyter notebook templates that were used to develop the paper figures. They provided expert

feedback on the performance and behavior of the `rubin_scheduler` and MAF metrics.

M.J. contributed significantly to the initial discussions about the implementation of the incremental templates to the LSST observing strategy cadence simulation as well as contributing to discussions involving early results with previous LSST survey strategy simulations.

W.C.F. and S.R.C. contributed to the development of the MAF metrics. They also provided feedback on the overall manuscript.
















J.K.P. provided guidance on the process of incremental template generation in the Rubin data management system and the incremental template generation strategy involved in Rubin DP0.2/DESC DC2. He also provided Figure 5.

S.R.M. and G.F. were involved in the early discussions that launched the analysis described in this work. They also provided feedback on the overall manuscript.

L.J.McG. provided a comparison of results at the healpixel and visit size scales.

B.T.B., C.O.C, S.G, H.H.H, and C.O. provided feedback on the overall manuscript.

ORCID iDs

James E. Robinson  <https://orcid.org/0000-0002-2121-6375>
 Megan E. Schwamb  <https://orcid.org/0000-0003-4365-1455>
 R. Lynne Jones  <https://orcid.org/0000-0001-5916-0031>
 Mario Jurić  <https://orcid.org/0000-0003-1996-9252>
 Peter Yoachim  <https://orcid.org/0000-0003-2874-6464>
 Bryce T. Bolin  <https://orcid.org/0000-0002-4950-6323>
 Colin O. Chandler  <https://orcid.org/0000-0001-7335-1715>
 Steven R. Chesley  <https://orcid.org/0000-0003-3240-6497>
 Grigori Fedorets  <https://orcid.org/0000-0002-8418-4809>
 Wesley C. Fraser  <https://orcid.org/0000-0001-6680-6558>
 Sarah Greenstreet  <https://orcid.org/0000-0002-4439-1539>
 Henry H. Hsieh  <https://orcid.org/0000-0001-7225-9271>
 Lauren J. McGinley  <https://orcid.org/0009-0007-5077-0475>
 Stephanie R. Merritt  <https://orcid.org/0000-0001-5930-2829>
 Cyrielle Opatom  <https://orcid.org/0000-0002-9298-7484>

References

Bellm, E., Blum, R., Graham, M., et al. 2020, Plans and Policies for LSST Alert Distribution LDM-612, STScI, <https://ldm-612.lsst.io/>
 Andreoni, I., Coughlin, M. W., Almualla, M., et al. 2022, *ApJS*, **258**, 5
 Bianco, F. & The SCOC 2022, Survey Cadence Optimization Committee's Phase 2 Recommendations PSTN-055, STScI, <https://pstn-055.lsst.io/>
 Bianco, F. & The SCOC 2024, Survey Cadence Optimization Committee's Phase 3 Recommendations PSTN-056, STScI, <https://pstn-056.lsst.io/>
 Bianco, F. B., Drout, M. R., Graham, M. L., et al. 2019, *PASP*, **131**, 068002
 Bianco, F. B., Ivezić, Ž., Jones, R. L., et al. 2022, *ApJS*, **258**, 1
 Bolin, B. T., Ahumada, T., van Dokkum, P., et al. 2022, *MNRAS*, **517**, L49
 Bolin, B. T., Fremling, C., Holt, T. R., et al. 2020, *ApJL*, **900**, L45
 Bolin, B. T., Masci, F. J., Coughlin, M. W., et al. 2025, *Icar*, **425**, 116333
 Bosch, J., Armstrong, R., Bickerton, S., et al. 2018, *PASJ*, **70**, S5
 Connolly, A. J., Angeli, G. Z., Chandrasekharan, S., et al. 2014, *Proc. SPIE*, **9150**, 915014
 Delgado, F., & Reuter, M. A. 2016, *Proc. SPIE*, **9910**, 991013
 Delgado, F., Saha, A., Chandrasekharan, S., et al. 2014, *Proc. SPIE*, **9150**, 915015

Fedorets, G., Micheli, M., Jedicke, R., et al. 2020, *AJ*, **160**, 277
 Feigelson, E. D., Bianco, F. B., & Bonito, R. 2023, *ApJS*, **268**, 11
 Gizis, J. E., Yoachim, P., Jones, R. L., Hilligoss, D., & Ji, J. 2022, *ApJS*, **263**, 23
 Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
 Graham, M. L., Bellm, E. C., Slater, C. T., Guy, L. P. & the DM System Science Team 2020, Options for Alert Production in LSST Operations Year 1 DMTN-107, STScI, <https://dmtn-107.lsst.io/>
 Guy, L. P., Bellm, E., Blum, B., et al. 2023, Rubin Observatory Plans for an Early Science Program RTN-011, STScI, <https://rtn-011.lsst.io/>
 Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
 Heinze, A., Eggl, S., Juric, M., et al. 2022, AAS DPS Meeting, **54**, 504.04
 Holman, M. J., Payne, M. J., Blankley, P., Janssen, R., & Kuindersma, S. 2018, *AJ*, **156**, 135
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111
 Ivezić, Ž. & The LSST Science Collaboration 2013, LSST Science Requirements Document LPM-17, STScI, <http://ls.st/LPM-17>
 Ivezić, Ž. & The SCOC 2021, Survey Cadence Optimization Committee's Phase 1 Recommendation PSTN-053, STScI, <https://pstn-053.lsst.io/>
 Jones, R. L., Slater, C. T., Moeyens, J., et al. 2018, *Icar*, **303**, 181
 Jones, R. L., Yoachim, P., Chandrasekharan, S., et al. 2014, *Proc. SPIE*, **9149**, 91490B
 Jones, R. L., Yoachim, P., Ivezić, Z., Neilsen, E. H., & Ribeiro, T. 2020, Survey Strategy and Cadence Choices for the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST), v1.2, Zenodo, doi:10.5281/zenodo.4048838
 Jurić, M., Axelrod, T., Becker, A., et al. 2023, Data Products Definition Document LSE-163, STScI, <https://lse-163.lsst.io/>
 Jurić, M., Eggl, S., Moeyens, J., & Jones, L. 2020, Proposed Modifications to Solar System Processing and Data Products DMTN-087, STScI, <https://dmtn-087.lsst.io/>
 Kubica, J., Denneau, L., Grav, T., et al. 2007, *Icar*, **189**, 151
 Lochner, M., Scolnic, D. M., Awan, H., et al. 2018, arXiv:1812.00515
 LSST Collaboration, Marshall, P., Anguita, T., et al. 2017, LSST Science Collaborations Observing Strategy White Paper: "Science-driven Optimization of the LSST Observing Strategy", v1.0, Zenodo, doi:10.5281/zenodo.842713
 LSST Dark Energy Science Collaboration, Abolfathi, B., Alonso, D., et al. 2021, *ApJS*, **253**, 31
 LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv:0912.0201
 LSST Science Collaboration, Marshall, P., Anguita, T., et al. 2017, arXiv:1708.04058
 McCully, C., Crawford, S., Kovacs, G., et al. 2018, astropy/astrocrappy: v1.0.5 Zenodo Release, Zenodo, doi:10.5281/zenodo.1482019
 McKinney, W. 2010, Proc. of the 9th Python in Science Conf., ed. S. van der Walt & J. Millman, (SciPy), 56
 Myers, J., Jones, L., & Axelrod, T. 2013, Moving Object Pipeline System Design LDM-156, STScI, <https://docushare.lsst.org/docushare/dsweb/Get/LDM-156/LDM-156.pdf>
 Naghib, E., Yoachim, P., Vanderbei, R. J., Connolly, A. J., & Jones, R. L. 2019, *AJ*, **157**, 151
 Rhoads, J. E. 2000, *PASP*, **112**, 703
 Schwamb, M. E., Jones, R. L., Chesley, S. R., et al. 2018a, arXiv:1802.01783
 Schwamb, M. E., Jones, R. L., Yoachim, P., et al. 2023, *ApJS*, **266**, 22
 Schwamb, M. E., Jurić, M., Bolin, B. T., et al. 2021, *RNAAS*, **5**, 143
 Schwamb, M. E., Volk, K., Wen, H., et al. 2018b, arXiv:1812.01149
 Shamir, L. 2005, *AN*, **326**, 428
 Swinbank, J., Axelrod, T., Becker, A., et al. 2020, Data Management Science Pipelines Design LDM-151, STScI, <https://ldm-151.lsst.io/>
 van Dokkum, P. G. 2001, *PASP*, **113**, 1420
 Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
 Yoachim, P. 2024, Lsst-Sims/sims_featureScheduler_runs4.0: Initial Release, v1.0, Zenodo, doi:10.5281/zenodo.13840868
 Yoachim, P., Coughlin, M., Angeli, G. Z., et al. 2016, *Proc. SPIE*, **9910**, 99101A
 Zonca, A., Singer, L., Lenz, D., et al. 2019, *JOSS*, **4**, 1298