LSST Survey Strategy in the Galactic Plane and Magellanic Clouds

R. A. Street[1,20], X. Li[2], S. Khakpash[2,3], E. Bellm[4], L. Girardi[5], L. Jones[4], N. S. Abrams[6], Y. Tsapras[7],
M. P. G. Hundertmark[7], E. Bachelet[8], P. Gandhi[9], P. Szkody[10], W. I. Clarkson[11], R. Szabó[12,13,14], L. Prisinzano[15],
R. Bonito[15], D. A. H. Buckley[16,17,18], J. P. Marais[18], and R. Di Stefano[19]

[1] Las Cumbres Observatory, 6740 Cortona Drive, Suite 102, Goleta, CA 93117, USA
[2] Department of Physics & Astronomy, University of Delaware, 217 Sharp Lab, Newark, DE 19716, USA
[3] Rutgers, The State University of New Jersey, Department of Physics & Astronomy, 136 Frelinghuysen Road, Piscataway, NJ 08854, USA
[4] DiRAC Institute, Department of Astronomy, University of Washington, 3910 15th Avenue NE, Seattle, WA 98195, USA
[5] Osservatorio Astronomico di Padova, INAF, Vicolo dell'Osservatorio 5, I-35122 Padova, Italy
[6] Department of Astronomy, 501 Campbell Hall #3411, University of California at Berkeley, Berkeley, CA 94720-3411, USA
[7] Astronomisches Rechen-Institut, Mönchhofstr. 12-14, D-69120 Heidelberg, Germany
[8] IPAC, Mail Code 100-22, Caltech, 1200 E. California Boulevard, Pasadena, CA 91125, USA
[9] University of Southampton, B46, West Highfield Campus, University Road, Southampton, SO17 1BJ, UK
[10] Department of Astronomy, University of Washington, Box 351580, Seattle, WA 98195-1700, USA
[11] Department of Natural Sciences, University of Michigan–Dearborn, 4901 Evergreen Road, Dearborn, MI 48128, USA
[12] Konkoly Observatory, CSFK, MTA Centre of Excellence, Budapest, Konkoly Thege Miklós út 15-17, H-1121 Hungary
[13] MTA CSFK Lendület Near-Field Cosmology Research Group, H-1121 Budapest, Konkoly Thege Miklós út 15-17, Hungary
[14] ELTE Eötvös Loránd University, Institute of Physics, 1117, Pázmány Péter sétány 1/A, Budapest, Hungary
[15] INAF—Osservatorio Astronomico di Palermo, Piazza del Parlamento 1, I-90134 Palermo, Italy
[16] South African Astronomical Observatory, P.O. Box 9, Observatory Road, 7935 Observatory, Cape Town, South Africa
[17] Department of Astronomy, University of Cape Town, Private Bag X3, Rondebosch 7701, South Africa
[18] Department of Physics, University of the Free State, P.O. Box 339, Bloemfontein 9300, South Africa
[19] Harvard–Smithsonian Center for Astrophysics, 60 Garden Street, Office P-227, Cambridge, MA 02138, USA

Abstract

Galactic science encompasses a wide range of subjects in the study of the Milky Way and Magellanic Clouds, from young stellar objects to X-ray binaries. Mapping these populations, and exploring transient phenomena within them, are among the primary science goals of the Vera C. Rubin Observatory's Legacy Survey of Space and Time. While early versions of the survey strategy dedicated relatively few visits to the Galactic Plane region, more recent strategies under consideration envision a higher cadence within selected regions of high scientific interest. The range of galactic science presents a challenge in evaluating which strategies deliver the highest scientific returns. Here we present metrics designed to evaluate Rubin survey strategy simulations, based on the cadence of observations they deliver within regions of interest to different topics in galactic science, using variability categories defined by timescale. We also compare the fractions of exposures obtained in each filter with those recommended for the different science goals. We find that the `baseline_v2.x` simulations deliver observations of the high-priority regions at sufficiently high cadence to reliably detect variability on timescales >10 days or more. Follow-up observations may be necessary to properly characterize variability, especially transients, on shorter timescales. Combining the regions of interest for all the science cases considered, we identify those areas of the Galactic Plane and Magellanic Clouds of highest priority. We recommend that these refined survey footprints be used in future simulations to explore rolling cadence scenarios, and to optimize the sequence of observations in different bandpasses.

Unified Astronomy Thesaurus concepts: Galactic and extragalactic astronomy (563); Milky Way Galaxy (1054); Magellanic Clouds (990); Galactic bulge (2041); Optical astronomy (1776); Optical observation (1169); Astronomical methods (1043); Time domain astronomy (2109)

1. Introduction

The National Science Foundation's Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) will cover the entire southern sky visible from the telescope's site at Cerro Pachón, Chile, but not to uniform limiting magnitude or temporal cadence. Figure 1 illustrates the number of visits realized to different regions of the sky, comparing the `baseline_v1.5` and `baseline_v2.0` iterations of the main Wide–Fast–Deep (WFD) survey strategy. Regions at lower ($|b| < 15°$) galactic latitude and extreme $\gtrsim 60°$ decl. received significantly fewer visits than the rest of the sky over the 10 yr span of LSST in `baseline_v1.5`. The distribution of visits was adjusted in `baseline_v2.0` to include the Magellanic Clouds and central Galactic Bulge, though much of the Galactic Plane still receives relatively few exposures.

This survey strategy places severe limitations on the science that the survey will realize, particularly for stars and planets in the Milky Way. Eleven of the 46 White Papers on the LSST survey cadence submitted in 2018[21] argued for increasing the number of visits to the Galactic Plane and Magellanic Clouds. These papers identified a vast array of scientific inquiries that

---

[20] Corresponding author.

---

[21] https://www.lsst.org/submitted-whitepaper-2018
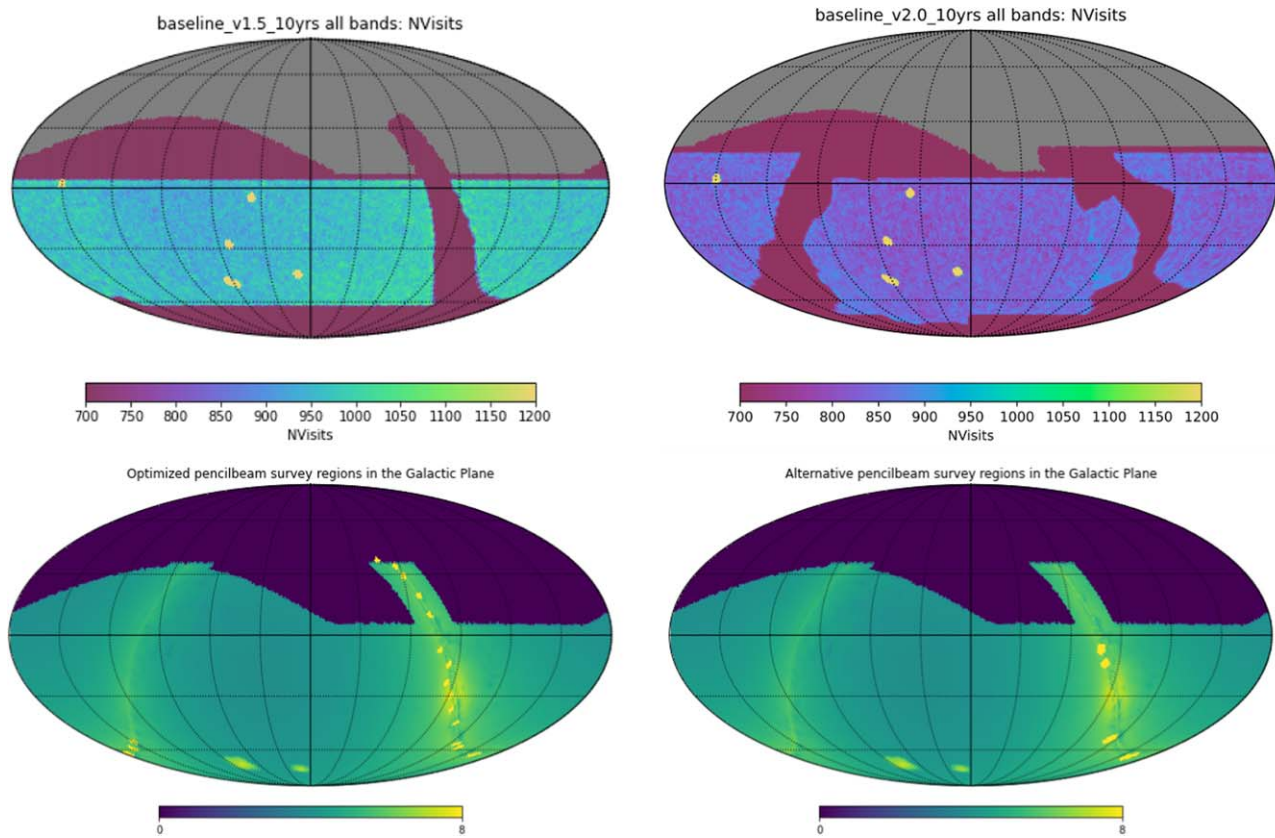
Figure 1. Top: number of visits per sky pointing over all years and filters, derived from the simulated v1.5 (left) and v2.0 (right) survey strategies. Bottom: sky maps with a normalized color gradient (values 0–1) indicating the $\log_{10}$ (density of stars) based on TRILEGAL models, overlaid with the two sets of Galactic Plane pencil-beam fields represented in yellow (values 1+; see Section 2.1.1). Set 1 is plotted on the left, set 2 on the right.

could be revolutionized by systematic observations from the Rubin Observatory, including:

1. Exploring the atmospheric properties of time-variable brown dwarfs and extending the known population of transiting planets and white dwarfs (Lund et al. 2018).
2. Comparing pre-main-sequence stars in star-forming regions (SFRs) with a range of metallicities down to the lowest-mass stars; using stellar variability and photometric properties to determine the ages and distances of open and globular clusters, and so informing our understanding of the evolution of the variables within those clusters (Prisinzano & Magrini 2018).
3. Using the photometric variability of young stellar objects to characterize mass accretion from circumstellar disks as well as magnetic cycles and flares (Bonito & Hartigan 2018).
4. Developing a more complete view of the population of cataclysmic variables (CVs) and AM CVn stars, detected from their outbursts and different accretion states, to be compared with close binary evolution models and to track the accretion behavior of white dwarfs (Olsen et al. 2018).
5. Detecting and monitoring outbursts from X-ray binaries that will reveal the population of neutron stars and black holes, determining whether a mass gap exists between them, establishing whether supernova explosions give large black hole kicks, and identifying them as a potential source of gravitational-wave events (Strader et al. 2018).
6. Producing extensive catalogs of stellar flares that will enable them to be studied as a function of stellar

properties, shedding light on the star's internal structure and processes; obtaining pulsations of white dwarfs to trace their cooling curves, and exploring pulsations as a function of metallicity; measuring the ages and metallicities of stars in the Galactic Bulge and tracing the formation history of the region (Bono et al. 2018; Gonzalez et al. 2018).

7. Using microlensing by stars and planets to reveal the distributions of these populations in different stellar environments throughout the galaxy, while also discovering isolated black holes and enabling the mass function of these objects to be fully described (Street et al. 2018b).
8. Using RR Lyrae stars, Cepheids, SX Phoenicis, delta Scuti stars, and long-period variables (LPVs) within resolved stellar populations to trace different stellar generations over the large spatial extension and magnitude depth allowed by the LSST; measuring their distances using variable stars of different types/parent stellar populations will map their 3D structures, and detect tidal streams (Clementini et al. 2018).
9. Detecting recurrent novae and determining whether they are progenitors of supernovae Type Ia (Strader et al. 2018).

Several papers highlighted LSST's ability to map the unique extragalactic stellar populations of the Magellanic Clouds, down to faint limiting magnitudes, if they were included in the survey footprint. LSST could therefore produce a complete catalog of variability for these populations, including detecting the microlensing signatures of planets (Olsen et al. 2018; Poleski & Mróz 2018; Street et al. 2018b). Using DECam

imaging of crowded fields as a proxy for LSST data, Suberlak et al. (2018) showed that a single 30 s exposure will detect, 5σ above the background sky, all objects brighter than $r_{lim} \leqslant 21.5$ mag in crowded regions, and probe to even fainter magnitudes, $r_{lim} = 24.5$ mag, in uncrowded areas of the sky.

Rubin observations of the Galactic Bulge can also enhance the scientific returns of other contemporary surveys, in particular the Nancy Grace Roman Space Telescope, which will provide very-high-cadence near-infrared time-series observations, with the goal of detecting planets and stars via their microlensing signatures, but only for a few periods of ~72 days each. Rubin observations could detect planets that would otherwise be missed by Roman, and enable the mass and distances of free-floating planets to be measured (Street et al. 2018a), if its survey of the Bulge region is coordinated with Roman, and realizes an adequate cadence of imaging.

Many of these science goals do not depend on achieving deep coadded images, in contrast to the WFD. Rather, their observational goals require repeated visits to the Galactic Plane and Magellanic Clouds in order to achieve time-series photometric measurements. In some cases, achieving the science goals would require a different cadence of observations in different filters, in light of the greater extinction compared with the WFD region.

LSST must meet its original scientific requirements (Ivezic & LSST Science Collaboration 2018), meaning that modifications to the survey strategy can only take up to a maximum of ~10% of the available on-sky time. These factors combine to argue that a distinct survey strategy should be considered for the Galactic Plane and Magellanic Clouds to maximize the scientific return within the survey requirements. The Rubin Observatory could run this survey in parallel with, or as part of, the WFD, during the same 10 yr period.

The Rubin Observatory designed the LSST Metric Analysis Framework (MAF; Jones 2021) to incorporate a range of metrics, including those contributed by the community, to evaluate the impacts of different survey strategies on different science cases. Many of the papers quoted above proposed metrics specific to each science topic, such as quantifying the numbers of a particular population that LSST will discover, or determining how well a given parameter will be measured. Detailed analyses of these science cases are presented in a number of parallel papers (e.g., Raiteri et al. 2022), and properly interpreting their outputs sometimes requires a nuanced, in-depth understanding of each science topic. The challenge for the Rubin Observatory is to find a strategy that optimizes the survey for a wide range of science, and to this end, the Rubin Survey Cadence Optimization Committee requested broader Figures of Merit that could be used as a guide when comparing different strategies for quite different science goals. The Rubin Observatory has recently conducted operational simulations (or "OpSims," as described in Bianco et al. 2022) of a wide range of alternative survey strategies, and invited community feedback on the scientific impacts of different configurations.

The goal of this paper is to review the observational needs of a range of different galactic science topics to identify areas of overlap and tension in the survey strategy, and to recommend scientifically motivated solutions. We present metrics designed to evaluate the essential observational parameters, so that the strengths and weaknesses of simulations of the different survey strategies may be evaluated, while representing a broad range of science goals.

In Section 2, we outline the essential components of any survey strategy, and describe our approach to capturing the requirements of galactic science for each component. In Section 3, we apply these metrics to compare the base-line_v2.0 strategy with previous iterations, and in Section 4 we evaluate the set of OpSim experiments conducted by Rubin in early 2022. In Section 5, we combine the regions of interest for different galactic science use cases to identify the highest-priority regions to survey, and summarize our key findings in Section 6.

2. Elements of Survey Strategy

In an ideal world, LSST would survey the entire visible sky with a regular cadence in all six filters. To do so would unfortunately reduce the achievable cadence within the WFD region to the point where the LSST Science Requirements are not met. More realistically, some trade-off is likely to be necessary between regular revisits to each telescope pointing on the sky (the cadence), the number and duration of exposures in different filters, and the on-sky area included in the survey footprint. For this reason, it is valuable to consider the survey footprint, cadence, and filter selection independently, so that metrics can evaluate the strengths of different strategies.

Studies of the populations of the Galactic Plane and Magellanic Clouds naturally follow the nonuniform spatial distributions of those populations. That is, galactic science needs observations in fairly well-defined regions at low ($|b| < 10°$) galactic latitude, with additional regions of interest centered on the Magellanic Clouds, star clusters, and SFRs. When the 2018 White Papers advocating for galactic science were compared, it was evident that many of them recommended observations of regions with a high degree of overlap. This led us to explore whether a common footprint could be defined, encompassing the regions of highest scientific priority.

The high extinction in much of the Galactic Plane tends to result in a preference for observations in redder passbands (g, r, i, z), particularly for time-series observations. y is not included in this set, because declines in the telescope and instrument throughput and detector sensitivity at these wavelengths produce a comparatively low signal-to-noise ratio in this filter. Nevertheless, observations in the u and y filters are still important for characterizing young star variability and spectral typing. The optimal total number and frequency of observations is therefore not uniform across all passbands.

By focusing on the observational requirements, rather than science-specific parameters, we can combine the needs of a diverse range of science in a concise set of metrics, which we describe in the following sections. For reference, the metrics are summarized in Table 1. All metric code has been integrated with the open-source Rubin MAF.[22]

2.1. Survey Footprint

In order to identify the regions of greatest scientific interest to galactic science, we first reviewed the White Papers described in Section 1. We compiled a list of the survey regions of interest from all papers citing galactic science as a motivation, together with the filter sets and cadences requested.

Table 1
Summary of Metrics Discussed in This Work

MAF Module	MAF Metric	Metric Output per HEALpix	Evaluated Metrics
galacticPlaneMetric	GalPlaneFootprintMetric	No. observations nObsPriority map_priority	%ofPriority(τ_{var}, map) %ofNObsPriority(τ_{var}, map)
	GalPlaneFilterMetric	$M_f(f, \alpha, \delta)$	% area $\geqslant 1.0$(map)
galplaneTimeSamplingMetrics	GalPlaneVisitIntervalsTimescaleMetric	VIM	%ofIdealVIM(τ_{var}, map)
	GalPlaneSeasonGapsTimescaleMetric	SVGM	%ofIdealSVGM(τ_{var}, map)

Note. The science-based priority maps are inputs to all these metrics, which are calculated using a HEALpix dataSlicer, meaning that the output quantities are calculated per HEALpix. Evaluated metric values are then derived by summing over HEALpix regions for different science cases and timescale categories of variability.

Table 2
Summary of the Larger-area Survey Regions Included in This Study

White Paper	Region	Gal. Long. ($l°$)	Gal. Lat. ($b°$)	Filters
Bono+ (deep survey)	Gal. Plane center	20–+20	−3–+3	*izy*
Bono+ (shallow survey)	Gal. Plane	20–+20	−15–+10	*ugriyz*
Gonzalez+	Gal. Plane center	−15–+15	−10–+10	*grizy*
Street+	Gal. Plane	−85.0–+85.0	−10.0–+10.0	*griz*
Prisinzano+, Bonito+	Gal. Plane/SFRs	−90.0–+90.0	−5.0–+5.0	*gri*
Poleski+, Street+, Clementini+	LMC	277.8–283.2	−35.2–30.6	*griz*
Street+, Lund+	Gal. Plane	−85.0–+85.0	−10.0–+10.0	*griz*
Poleski+, Street+, Clementini+	SMC	301.5–304.1	−45.1–43.6	*griz*
Street+(a)	Gal. Bulge	2.216	−3.14	*griz*
Clementini+	M54	5.60703	−14.08715	*gri*
Clementini+	Sculptor	287.5334	−83.1568	*gri*
Clementini+	Carina	260.1124	−22.2235	*gri*
Clementini+	Fornax	237.1038	−65.6515	*gri*
Clementini+	Phoenix	272.1591	−68.9494	*gri*
Clementini+	Antlia2	264.8955	11.2479	*gri*
Kharchenko+	Open clusters	Catalog	Catalog	
Baumgardt & Hilker	Globular clusters	Catalog	Catalog	

The sources are summarized in Table 2, including the selection of filters indicated to be most important to each science case. Many papers placed high value on *g*, *r*, *i*, *z* observations, particularly those concerned with targets in regions of high extinction or where a high cadence in a single filter is preferable, but we note that this should not be interpreted to mean that *no* data are required in *u* or *y*. Data in these filters are valued, but not always required at the same survey cadence. This list of regions was extended to include the curated catalogs of open clusters published in Kharchenko et al. (2013) and globular clusters described in Baumgardt & Hilker (2018), as well as SFRs (Zucker et al. 2020), which are not adequately represented by a simple function of N_{stars}/deg^2. In addition, we included a survey region optimized for the study of low-mass X-ray binaries (LMXBs) from a function that produces probabilistic samples of LMXB positions and distances based on a spatial distribution derived from the galactic mass model (Dehnen & Binney 1998; Grimm et al. 2002), following the method of Johnson et al. (2019; E. Bellm 2023, private communication).

Table 2 shows the extensive overlap between the recommended survey regions, and to first order the superset of these regions could be used as the footprint for a galactic science survey.

The regions of greatest interest closely follow the sky regions of highest stellar density, owing to the nature of the science concerned. This varies strongly, particularly across the Plane, and in general the regions of extreme extinction and lower stellar density are of lower (though not zero) scientific priority. We can therefore refine the survey footprint by prioritizing regions with higher stellar density.

It is valuable to present the footprint information in the form of a Hierarchical Equal Area isoLatitude Pixelisation (HEALpix; Gorski et al. 2005) map of the sky, where each HEALpix is assigned a numerical metric value representing the scientific priority of that HEALpix for galactic science. Any such metric is an unavoidably crude measure of true value, and does not diminish the importance of downweighted regions. Instead the purpose of this "priority map" is to indicate the regions where the majority of observations would ideally be directed in order to answer the needs of the science cases considered. This presents the information in a form that can then be combined with other metrics to evaluate how different LSST OpSims perform for these science cases, and can also be used within the LSST scheduler to determine the field priority in observation sequencing.

We combined the overlapping survey regions in Table 2 in the form of a HEALpix map with NSIDE = 64, resulting in HEALpix sizes of 55′.0. First, we used a map of stellar density in the *r* band derived from the TRILEGAL galactic model (Dal Tio et al. 2022), and identified the regions of highest interest by selecting sets of HEALpix with densities greater or equal to

60%, 70%, and 80% of the maximum stellar density over the whole sky. The pixels in these regions were assigned priority weightings of 0.8, 0.9, and 1.0, respectively. This map highlighted the Galactic Plane, Bulge, and Magellanic Clouds.

We then added to this map the other regions of special interest, such as SFRs, etc., which were initially assigned a priority weight of 1.0.

We opted to represent the bandpass requirements for different science cases in the following way. For each science use case, each of the six filters was assigned a weighting factor to represent the relative frequency with which observations in that filter should be obtained to meet that use case. Summed over all passbands, the weights add up to 1.0 for a given use case. For example, the filter weightings assigned for fields in the Galactic Plane were u: 0.05, g: 0.225, r: 0.225, i: 0.225, z: 0.225, and y: 0.05, capturing the lower frequencies of observations requested in the extreme red and blue bandpasses.

This enabled us to generate a HEALpix priority map for each filter by iterating over the science cases and assigning a priority for each HEALpix that is effectively the product of the HEALpix priority map and the filter weighting factor, summed over all the science cases considered. In this way, each science case effectively constitutes a "vote" toward the priority of the observations of a given HEALpix in a given filter.

2.1.1. Galactic Plane Pencil Beams

It was apparent from the first iteration of the priority maps that the many science interests in regions around the Galactic center combine to produce a strong weighting toward that part of the sky. However, surveying *only* this region at a moderate cadence neglects a key aspect of Galactic Plane science that the Rubin Observatory can uniquely excel in—namely, its ability to compare populations derived from different regions of the Milky Way.

A number of science drivers benefit from surveying a wider range of galactic longitudes.

1. Microlensing detection of black holes, stellar systems, and planetary systems: Surveying a range of galactic longitudes will discover faint or unseen populations along different lines of sight through the Galaxy. These lensing events will be caused by different populations in the spiral arms, disk, Bulge, or halo, depending on the line of sight, and comparing the rates and properties of microlensing events in different fields will produce insights into the population distributions and how they correlate with, e.g., galactic trends in metallicity.
2. Galactic structure: by surveying across galactic longitudes, LSST will characterize RR Lyrae and many other indicator variable types, which can be used to more accurately map structures such as the bar, as well as resolved stellar populations like globular clusters.
3. Detection of SFRs and stellar clusters: LSST coadded images will reach limiting magnitudes of >24.9 mag in all filters, enabling it to detect faint, low-mass populations of galactic clusters and SFRs. Open clusters and SFRs are distributed in a wide ($\sim20°$) band extending along the Galactic Plane.

That said, surveying the entire Milky Way Plane within the region $-85° \leqslant l \leqslant +85°$, $-10° \leqslant b \leqslant +10°$ at high cadence is likely to require more time than is available. The region encompasses approximately 354 distinct Rubin pointings,

Table 3
Summary of Two Sets of Galactic Plane Pencil-beam Fields

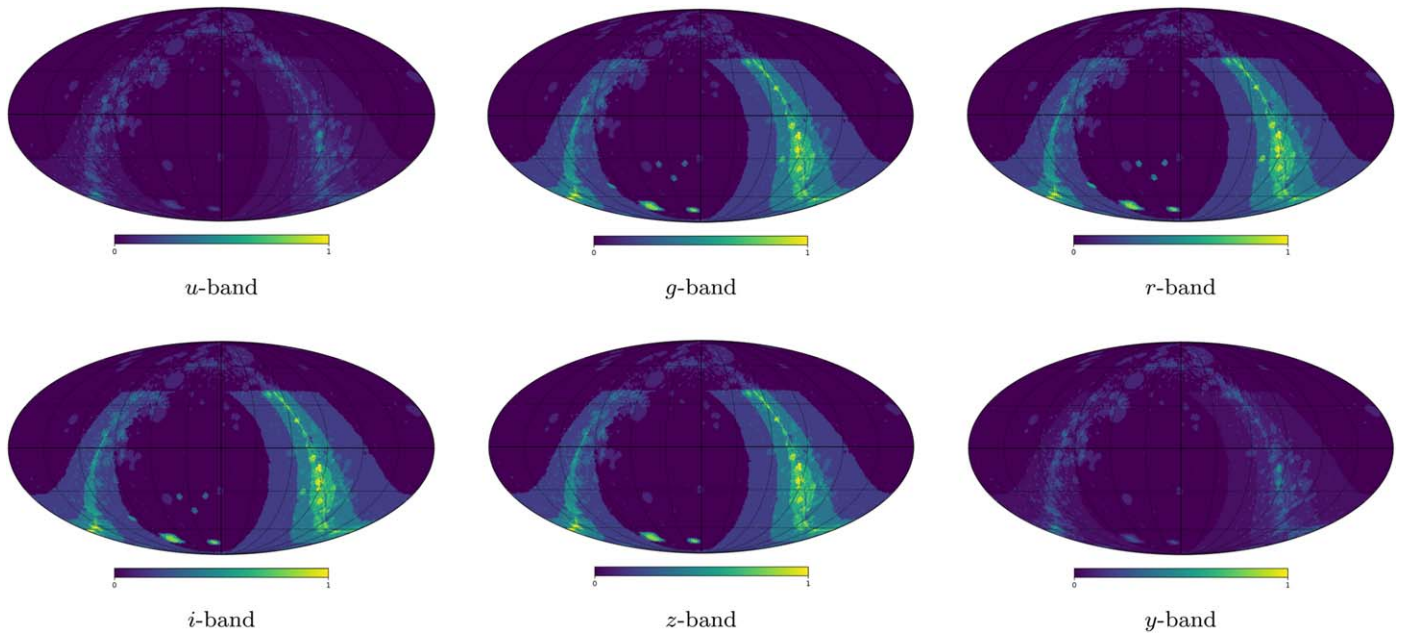| Field Set 1		Field Set 2	
l (deg)	b (deg)	l (deg)	b (deg)
280.0	0.0	306.56	−1.46
287.280702	0.0	331.09	−2.42
295.394737	−0.4167	18.51	−2.09
306.425439	−0.4167	26.60	−2.15
306.206140	−0.4167		
320.153509	−0.4167		
324.517544	−0.4167		
341.381579	−0.4167		
351.578947	−2.5000		
0.109649	−2.0833		
0.307018	−2.0833		
8.421053	−3.3333		
17.36842	−0.4167		
26.31579	−2.9167		
44.01316	−0.4167		
44.21053	−0.4167		
54.40789	0.0		
66.271930	−0.4167		
71.885965	0.0		
80.0	−5.0		

Note. Set 1 fields are single pointings of the Rubin Observatory with a radius of 1°.75, whereas the set 2 fields have a radius of 3°.91.

which (assuming typical visit and slew lengths of 42 s and 30 s, respectively) would take ~7 hr to image once per pointing. To survey this entire region at the nominal cadence of the WFD survey would take $\sim35\%$ of the available on-sky time per year. This easily exceeds the estimated 10% of time remaining after Rubin's Science Requirements are met.

In order to achieve the scientific goals above without compromising the Science Requirements, we modified the priority maps from Section 2.1 with a series of narrow "pencil-beam" fields distributed across the Galactic Plane (Table 3). Each field would represent a single Rubin pointing, and ~20 such pencil beams would be clustered around $b \sim 0°$, but distributed at regular intervals in galactic longitude.

To locate the pencil-beam fields, we began with 20 pointings distributed at regular intervals between $-85° \leqslant l \leqslant +85°$, and $b = 0°$, and then refined them to take into account the uneven extinction across the Galactic Plane. This was done using maps of the stellar density across the Plane derived from the TRILEGAL galactic model (Dal Tio et al. 2022). Each initial pencil-beam pointing was dithered by $\pm5°$ in l and b and optimized by selecting the location that maximized the stellar density for the field. Figure 1 (right panel) shows the optimized locations of the proposed fields.

The dithering around each survey pointing that is built into the Rubin Observatory scheduler raised a potential issue for the pencil-beam fields. Each survey pointing is designed to be dithered by up to 0°.7 (Jones 2021). While this expands the area of a spatially isolated field, and hence the total number of stars included, it also means that not all stars will be observed during every visit. For spatially contiguous survey regions, the dithered pointings will overlap a different area of the survey region, but for a set of spatially isolated fields like the pencil beams, the scientific benefits of including more stars could be overshadowed by the overall drop in cadence. To explore this

Figure 2. Sky maps where each HEALpix is weighted according to a priority metric for a broad combination of Galactic Plane science, presented for each filter, and including the field set 1 pencil beams.

potential issue, we also considered an alternative set of pencil-beam fields, which were designed to cover the same total spatial area as the original set, but within only four survey regions. The locations of these four regions were based on the EROS2 survey fields toward the Milky Way spiral arms (Moniez et al. 2017). Figure 1 allows the two sets of pencil-beam fields to be compared, and the field locations are given in Table 3. It should be noted that since the Galactic Bulge and Magellanic Clouds are considered to be separate survey regions and included independently of the pencil beams, they are *not* part of these field sets. A comparison of the two sets is discussed further in Section 5.

The pencil-beam fields were then included in the combined priority maps described above, with filter weightings of u: 0.1, g: 0.2, r: 0.2, i: 0.2, z: 0.2, and y: 0.1, to ensure that at least a minimum range of fields outside the Galactic Bulge region receive a comparable priority to the central Bulge region. All pencil-beam fields receive the same priority weighting per filter. The filter balance chosen reflects the preference for *griz* for enhanced time cadences in these fields, but allows for slightly more weight in u, y than the central Bulge, as many of the pencil beams have less well-characterized stellar populations and lower extinction. Although LSSTCam's quantum efficiency is lower in the u and y bands, data in these filters are still important to various scientific goals. u-band observations are important for measuring $\log(g)$ for white dwarfs, for example, and for distinguishing hydrogen and helium white dwarfs, as well as for determining stellar metallicities and as a measure of accretion activity in young stars (Bonito & Hartigan 2018; Bonito et al. 2023). This is particularly important for red clump stars in the Galactic Bulge (Johnson et al. 2020), despite the high extinction in that region resulting in lower signal-to-noise ratios. y-band data are needed for the detection of brown dwarfs, and to characterize instances of their variability due to cloud structures in their atmospheres (Street et al. 2018a; Tan & Showman 2019). The final priority maps with the set 1 pencil beams are shown in Figure 2.

The code used to generate these priority maps, as well as the maps themselves in FITS and PNG form, can be found in the TVS GitHub repository.[23]

2.1.2. Survey Footprint Metric

The footprint maps described above are designed to be *tunable*, in the sense that a larger or smaller survey area can be chosen by setting the HEALpix priority threshold to lower or higher values, respectively. This allows trade-offs between survey area and cadence to be explored in different simulations.

To evaluate this in a more quantified manner, the `GalPlaneFootprintMetric.py` calculates two metrics. It accepts one of the science priority maps and a HEALpix dataSlice as input parameters, and selects observations in each filter in turn, where the exposures reach a minimum 5σ depth of u: 22.7, g: 24.1, r: 23.7, i: 23.1, z: 22.2, and y: 21.4 mag. These limits were derived from the limiting magnitudes where LSST's catalog of detected sources is expected to be complete to $\geqslant 50\%$ in crowded fields (Suberlak et al. 2018). It returns both the number of observations per HEALpix matching these criteria, together with the number of observations multiplied by the priority of the HEALpix in question. It returns a null value for HEALpix lying outside the desired survey region of the given priority map for each science case, so a summary metric can be calculated simply by summing over all HEALpix.

The metric calculates the number and priority of useful observations of a given HEALpix within the scientifically desired survey region. In order to meaningfully interpret these values, we now need to consider what cadence of observation is needed to meet the science needs.

2.2. Variability Timescales and Survey Cadence

Time-domain astrophysics, including both transient and repeating variable phenomena, is the science driver behind much, though not all, of the galactic science use cases

---

[23] https://github.com/LSST-TVSSC/software_tools

considered here. Accurately capturing all nuances of every type of variability in MAF-style metrics is challenging, and it risks the possibility of the nuance being lost or misunderstood in the interpretation of the results.

We propose that this goal can be approached by considering the following broad categories of variability according to the typical timescales of the phenomena, e.g., τ_{var}: <10 days, 10–100 days, 100–365 days, and >365 days. These categories can include both transient and (quasi)periodic phenomena if τ_{var} is ascribed to either the duration of a transient event or the cycle period, respectively. Irregular phenomena may be considered to be a subclass of either transients (if they exhibit extended periods of quiescence between excited states) or quasiperiodic phenomena (if the variability is near continuous, τ_{var} can refer to the average time between high and low states). For current purposes, the critical aspect is how quickly or how often a variable target changes photometric state, since the survey would need to observe at a minimum of this frequency in order to reliably detect the variability. Examples of these categories include the following science use cases.

1. <10 days includes transits by exoplanets and white dwarfs, pulsations, stellar flares, and self-lensing by compact objects in binaries, short-period stellar binaries, pulsating stars (including RR Lyrae and Cepheids), some microlensing (e.g., by free-floating planets, bound planet anomalies, and brown dwarfs), variability in young stars, and short-term accretion variability.
2. 10–100 days includes most microlensing by galactic stellar-mass lenses, intermediate pulsation periods, short-term disk and magnetic instability in CVs, novae, and accretion in young stars.
3. 100–365 days includes microlensing by stellar remnants, long-timescale pulsations, medium-timescale disk instability, and mass transfer in CVs.
4. >365 days includes LPVs, e.g., Miras, microlensing by black holes, long-term disk instability, and mass transfer in CVs.

It should be noted that periodic objects can be detected from low-cadence data if the baseline is long enough, especially in the shortest-timescale categories (see, for example, Wiktorowicz et al. 2021). In terms of survey strategy, this means that the variability timescale does not directly set the minimum cadence requirement. For these science goals, a better metric is to analyze how well LSST will measure different periods, which is explored in parallel papers, including Geller et al. (2021), Di Criscienzo et al. (2023), and A. Prsa (2023, in preparation).

This work therefore focuses primarily on transient, irregular, and quasiperiodic variability. However, there is an argument in favor of using the following metrics for periodic objects, particularly within the first few years of LSST, to ensure that periodic objects are identified as soon as possible. In addition to promoting the early study of these variables, it will also improve the identification of transients, as they can be excluded from classification.

Although LSST will generate an alert whenever an object varies by more than 5σ from the template image flux, for most of the galactic science a single detection is insufficient to reliably identify variability. Based on previous follow-up experience (e.g., Tsapras et al. 2019), we require a minimum of five observations taken within τ_{var} to consider a variable
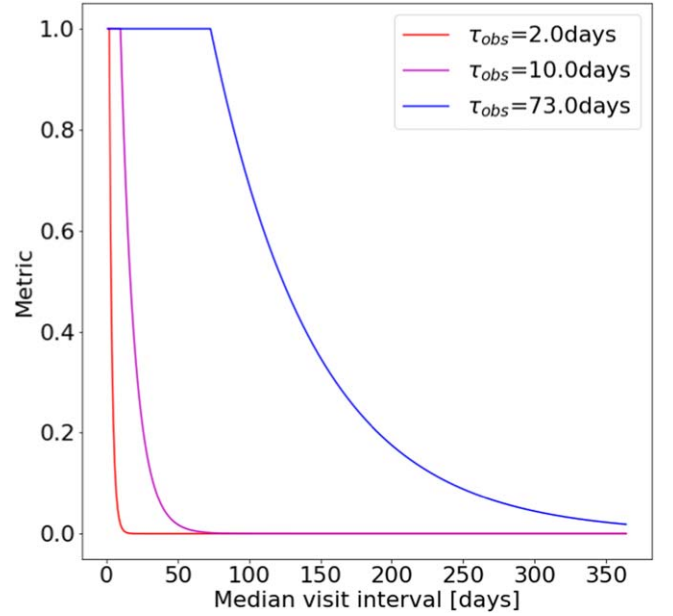
Figure 3. Merit function representing how the scientific value of sequential observations of time-variable objects decreases as the interval between observations exceeds the timescale characteristic of the object's variability.

target "detectable," in the sense of being able to accurately extract targets of interest from the LSST datastream. It should be noted that this neglects requirements for multifilter observations within a given timescale, as filter selection is considered with an independent metric described below.

Applying this requirement to transient phenomena, we derive a maximum useful observation interval, $\tau_{\mathrm{obs}} = \tau_{\mathrm{var}}/5$, for the categories above ({ } notation indicates where multiple categories are used to cover a span of timescales):

1. $\tau_{\mathrm{var}} < 10$ days, $\tau_{\mathrm{obs}} = 2$ days;
2. $10 < \tau_{\mathrm{var}} < 100$ days, $\tau_{\mathrm{obs}} = 2$–20 days; represented as $\tau_{\mathrm{obs}} = \{5, 11\}$ days;
3. $100 < \tau_{\mathrm{var}} < 365$ days, $\tau_{\mathrm{obs}} = 20$–73 days; represented as $\tau_{\mathrm{obs}} = \{20, 46.5\}$ days; and
4. $\tau_{\mathrm{var}} > 365$ days, $\tau_{\mathrm{obs}} = 73$ days.

Two values of τ_{obs} are considered for the τ_{var} categories between 10 and 365 days, as they encompass a large number of different phenomena of interest that span this timescale range. We therefore consider six representative categories of variability in what follows.

Typically, observations of time-domain variability, particularly transient phenomena, are most valuable when they occur at the intervals of τ_{obs}. However, they are usually still valuable even if they occur after this interval, though the value of later observations decreases with an increasing interval. We represent this as a visit interval metric (VIM), a decaying function of the interval between sequential visits (plotted in Figure 3), Δt and time t, where

$$\mathrm{VIM} = 1.0, \, t \leqslant \tau_{\mathrm{obs}}, \tag{1}$$

$$\mathrm{VIM} = e^{-K(\Delta t - \tau_{\mathrm{obs}})}, \, t > \tau_{\mathrm{obs}}, \, K = 1.0/\tau_{\mathrm{obs}}. \tag{2}$$

Even if observations are acquired at a regular cadence appropriate to the variability of a given target, there is a second factor impacting whether the target can be characterized from the resulting time series. With the exception of circumpolar objects, a given object is only visible to Rubin for up to a few

months each year. For variability categories where $\tau_{\rm obs}$ is shorter than the gap between seasons, $\Delta_{\rm gap}$, the season length is not the dominant factor in evaluating a given strategy, and a season visibility gap metric (SVGM) should return 1.0. But objects with longer-variability timescales can be difficult to characterize if the gaps between seasons are comparable to the $\tau_{\rm var}$, so we apply a similar decay function:

$$\text{SVGM} = 1.0, \ t \leqslant \tau_{\rm obs}, \tag{3}$$

$$\text{SVGM} = e^{-K(\Delta_{\rm gap} - \tau_{\rm obs})}, \ \Delta_{\rm gap} > \tau_{\rm obs}, \ K = 1.0/\tau_{\rm obs}. \tag{4}$$

Consideration of these two components can be used to evaluate how well the time series from a given survey strategy will allow variability in these timescale categories to be detected.

2.3. Filter Selection

The maps of the desired survey regions described in Section 2.1 were designed to take into account the relative priority placed on observations in different filters by different science cases, by assigning a weighting between 0 and 1 to each filter for each science case, and then multiplying by the priority of the pointing for each HEALpix. Relative weights were then summed over all science cases considered to create a priority map in each filter. These filter-specific priority maps can be used in a metric to evaluate how well a given OpSim meets the filter requirements of the science in the following way.

For each survey strategy simulation, the summed exposure time spent per pointing per filter, f, can be calculated as a function of spatial position for each HEALpix as a function of the total exposure time dedicated to that HEALpix, $(R_{\rm exp})$:

$$R_{\rm exp}(f, \alpha, \delta) = \frac{\sum_0^{N\,{\rm exp}} t_{\rm exp}(f, \alpha, \delta)}{\sum_0^{N\,{\rm exp}} t_{\rm exp}(\alpha, \delta)}, \tag{5}$$

where f indicates the filter bandpass, α, δ denote the R.A. and decl. of the HEALpix, $t_{\rm exp}$ refers to the exposure time in seconds, and $N_{\rm exp}$ gives the number of exposures in each filter per HEALpix. In the optimum survey strategy, these proportions per HEALpix should correspond to the relative priorities, $p(f, \alpha, \delta)$, of that HEALpix in the priority maps for each filter. This is estimated as a fraction of the total priority for that pixel summed over all filters:

$$F_W(f, \alpha, \delta) = \frac{p(f, \alpha, \delta)}{\sum_f p(f, \alpha, \delta)}. \tag{6}$$

The metric therefore outputs $M_f(f, \alpha, \delta) = R_{\rm exp}(f, \alpha, \delta)/F_W(f, \alpha, \delta)$, and this quantity is evaluated by calculating the percentage of the desired survey HEALpix where $M_f \geqslant 1.0$.

It is informative to evaluate the above metrics for footprint, survey cadence, and filter selection individually, for the information they provide regarding the strengths and weaknesses of a given survey strategy for different types of variability. We explored the option of producing a single Figure of Merit that calculated the product of these factors to see if it provided a useful guide to survey strategy choices. However, in practice, we found that this inevitably obscures significant nuance for the many different science cases included, so we proceeded with a careful evaluation of the individual component metrics.

3. Evaluating the `baseline_v2.0` Strategy

Following community feedback to Jones (2021), the then-nominal survey footprint for the WFD (`baseline_v1.5`) was modified to include a region around the Galactic Bulge (`baseline_v2.0`). It is instructive to compare these survey footprints with the galactic science survey footprint from Section 2.1, to quantify how well they would serve the science drivers considered here. We note that a `baseline_v2.1` simulation was also produced, which included the Virgo cluster in the low-dust WFD region and added a requirement to acquire r- and i-band images in conditions of good seeing (FWMH$_{\rm effective} < 0\rlap{.}''8$) each year. Although all simulations were analyzed, in practice we found that these changes have minimal impact on the following discussion. Since `baseline_v2.0` continues to be used widely as the canonical reference at the time of writing, we base our comparisons on it. It should also be noted that each visit to an on-sky pointing consisted of 1×30 s exposures in `baseline_v1.5` compared with 2×15 s for g, r, i, z, y filter exposures in the `_v2.x` simulations. This means that there are $\sim 9\%$ more visits available in `_v1.5` compared with later implementations, due to the additional overheads.

For each LSST survey strategy simulation, we identified the region surveyed as the area including all the HEALpix to receive a minimum number of total visits during the 10 yr survey. Observations in all filters are included, but we required that the limiting magnitude was equal to or exceeded 21.5 mag. This represents the faintest object expected to be detected at 5σ in a single LSST exposure, in any filter, of nominal duration of 30 s, in crowded Galactic Plane fields with moderate extinction (Suberlak et al. 2018). While this is a conservative limit in some bandpasses, it was adopted in all filters to simplify the code, because the effect of this constraint is simply to exclude short exposures in a few OpSims designed to explore this option. Thresholds for the minimum number of visits were estimated by considering the maximum useful interval between observations, $\tau_{\rm obs}$, described in Section 2.2. We approximated the average seasonal visibility of a given HEALpix from the Rubin Observatory to be 6.5 months, taking into account the requirement for two visits per pointing per night. From this, we can estimate the expected number of visits to a given HEALpix over the course of the 10 yr survey if the survey achieves the interval $\tau_{\rm obs}$, for each of the four variability categories.

The `GalPlaneFootprintMetric` was used to calculate the percentage of HEALpix in the desired survey footprint to receive the minimum number of visits in each filter. We considered both the combined survey region, as well as the survey regions for each science case separately, to gain better insight into the science that most benefits from each survey strategy. Figures 4 and 5 summarize the percentages of the desired regions to receive the minimum number of visits needed to characterize the six categories of variability timescale (see Section 2.2), for the `baseline_v1.5` and `baseline_v2.0` simulations, respectively.

Figure 4 shows clearly that no region received sufficiently high cadence in `baseline_v1.5` to characterize the most rapidly variable targets ($\tau_{\rm obs} = 2$ days/$\tau_{\rm var} < 10$ days) from Rubin data alone. Notably, the key regions of the Galactic Bulge and Magellanic Clouds did not receive sufficient observations to characterize even intermediate-timescale ($\tau_{\rm obs} = 5$, 11 days) variables either, although variables in categories with longer timescales are detectable in all regions.
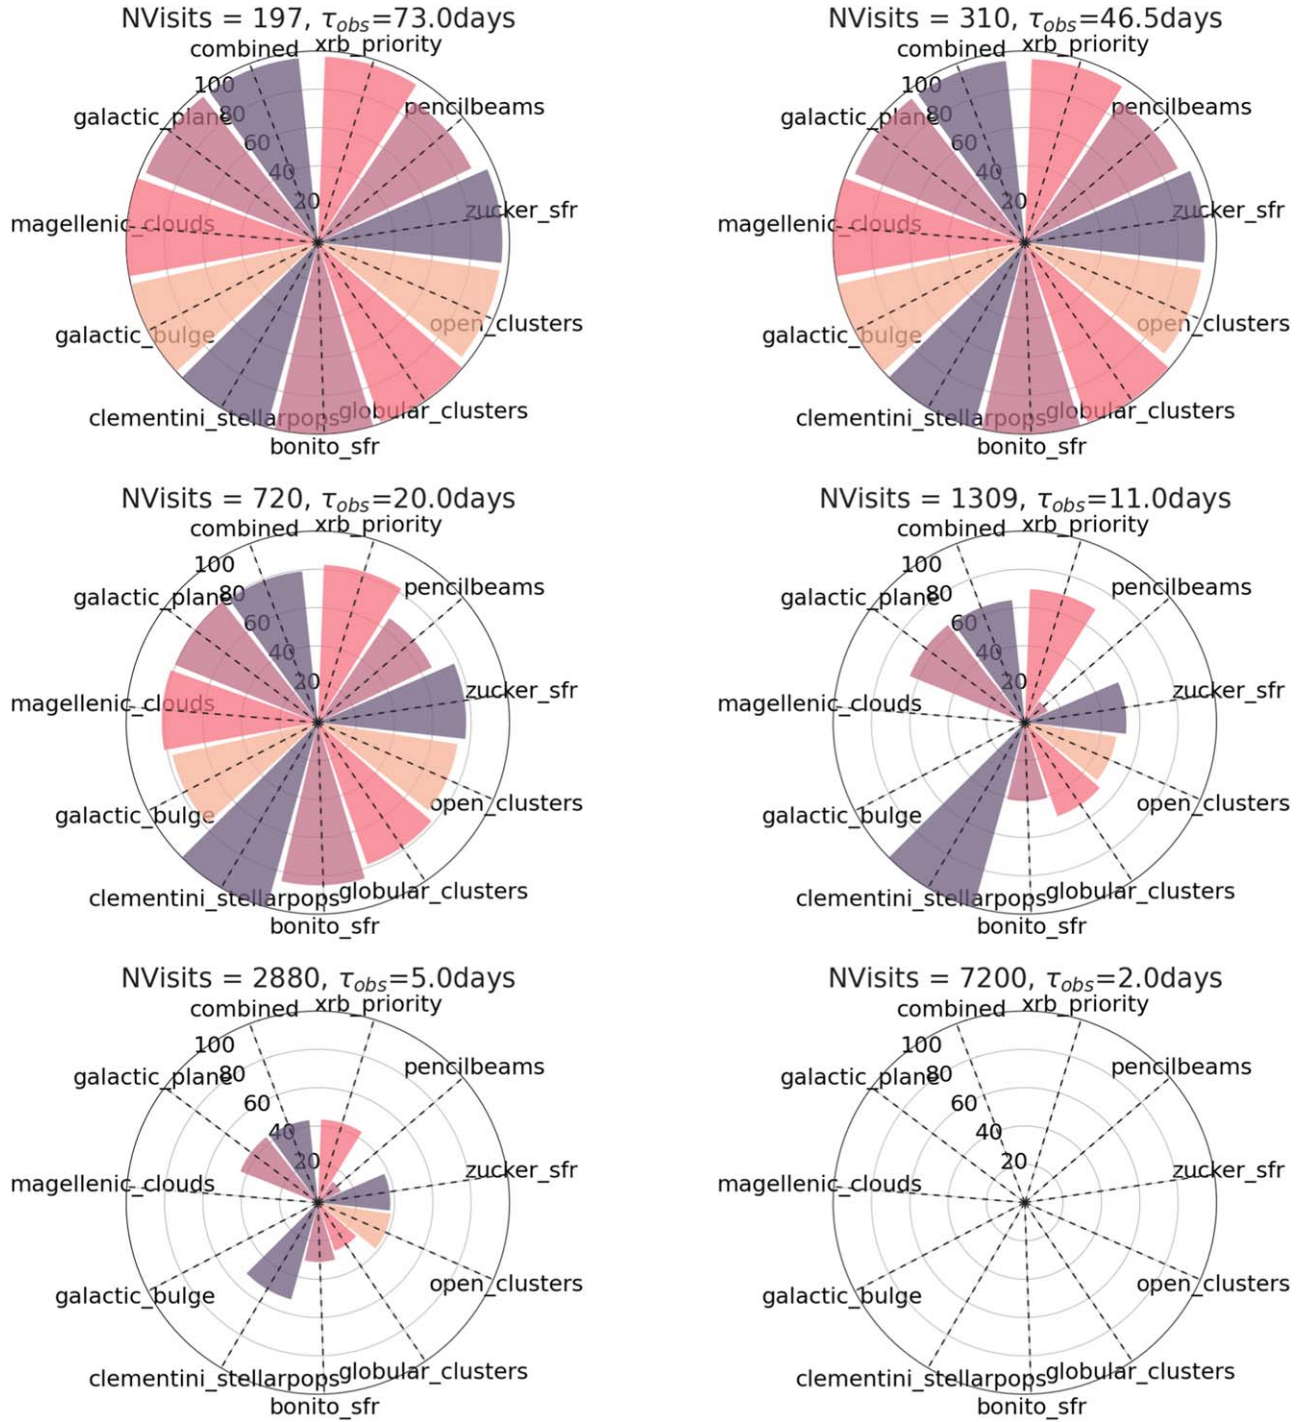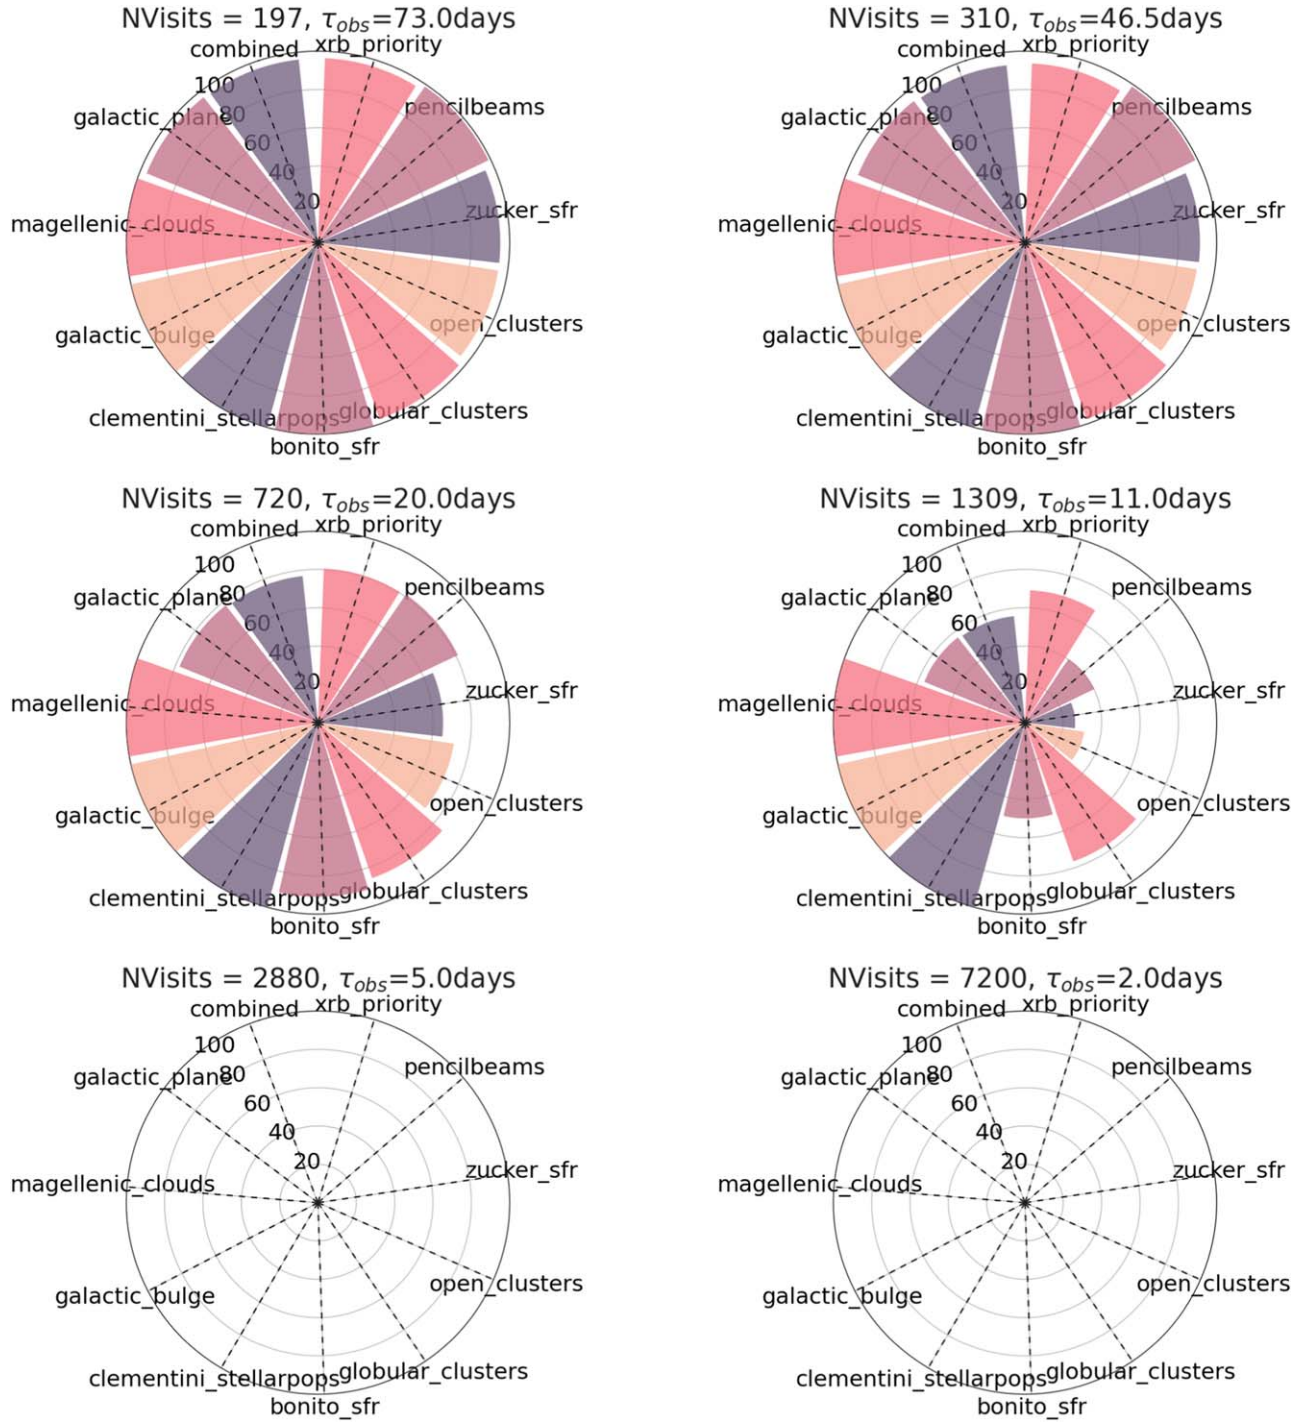
Figure 4. The percentages of the desired survey regions to receive the indicated number of visits over the 10 yr survey lifetime in the `baseline_v1.5` simulation (%ofPriority metric), plotted for the four categories of variability timescale.

Strikingly, ∼40% of the combined map of desired regions received sufficient visits to detect variability with $\tau_{\mathrm{obs}} > 5$ days in `baseline_v1.5`, while no region did in `baseline_v2.0` (Figure 5). This is because `baseline_v1.5` dedicated more visits to areas in the Galactic Plane with $5 < |b| < 15°$, whereas `baseline_v2.0` concentrated on the central Galactic Plane and Magellanic Clouds. This means that intermediate-timescale variables ($\tau_{\mathrm{obs}} \geqslant 11$ days) are very well recovered from those regions in `baseline_v2.0` data. The results are similar for variables in the intermediate categories for most other regions except globular clusters,

which see an improvement of ∼25%. The lack of high-cadence data in any Galactic Plane region significantly curtails the detected population of microlensing events, and other transient phenomena, including X-ray binary/CV outbursts.

This illustrates one possible trade-off between sky area covered and time cadence achieved, and while the additional visits to the Galactic Bulge and Magellanic Clouds enhance a range of science, we argue that further optimization is possible. In particular, we note that a great deal of high-priority science in the short τ_{obs} categories is lost when no areas receive visits at high cadence. This is illustrated in Figure 6, which presents the

Figure 5. The percentages of the desired survey regions to receive the indicated number of visits over the 10 yr survey lifetime in the `baseline_v2.0` simulation, plotted for the four categories of variability timescale.

regions of the combined HEALpix science priority map to receive $\geqslant 1309$ visits ($\tau_{obs} > 11$ days) in the `baseline_v1.5` and `_v2.0` simulations. While `baseline_v2.0` allocates more visits to high-priority regions, it still includes more visits in relatively low-priority fields at $|b| > 10°$, instead of the rest of the Galactic Plane at galactic longitudes away from the Bulge. The first option for the `GalacticPlaneMetric` "% ofPriority" sums the science priority of all HEALpix within a given science region to receive sufficient numbers of visits to meet the sampling requirement for each τ_{obs} category. This value is presented as a percentage of the summed priority of all

HEALpix within the region in question, giving a measure of whether a given observing strategy is focused on the highest-priority regions. We use this to compare `baseline_v1.5` and `_v2.0` in Figure 7, for $\tau_{obs} = 11$ days. This shows that the revised `baseline_v2.0` strategy benefits several science cases, with more of their high-priority regions being monitored, notably the Bulge, Magellanic Clouds, globular clusters, X-ray binaries, and the Bonito set of SFRs, and an increase of over 26% in the monitoring of the pencil-beam fields.

In contrast, the percentage of HEALpix from the Galactic Plane region monitored at this cadence actually decreases
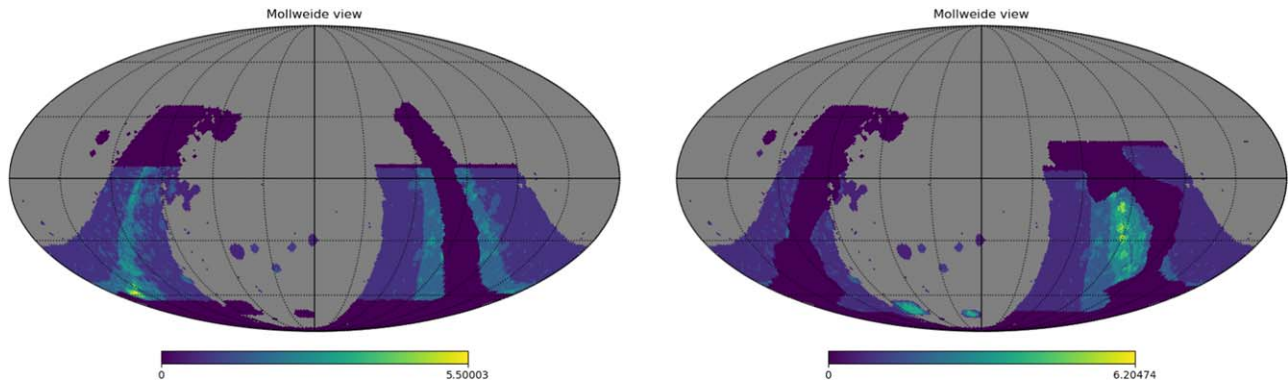
Figure 6. Maps of the combined science priority of each HEALpix region to receive at least 1309 visits during LSST (the average cadence adequate to detect variability on intermediate timescales of ∼11 days). The dark areas indicate which scientifically desirable regions are insufficiently observed. Left: `baseline_v1.5`. Right: `baseline_v2.0`.
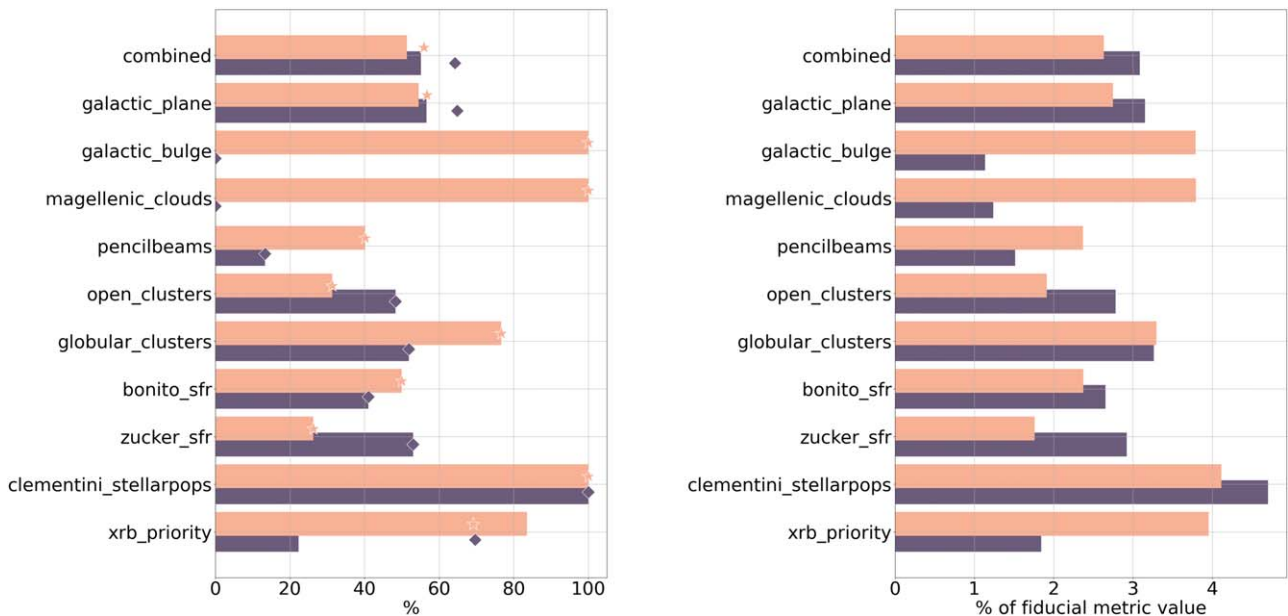
Figure 7. Galactic Plane Figures of Merit used to compare `baseline_v1.5` (dark purple, diamonds) and `_v2.0` (salmon pink, stars). Left: summed science priorities of the HEALpix within a given science region to receive adequate sampling in the $\tau_{obs} = 11$ days ($\tau_{obs} = 5$ days) category (bars), compared with the number of HEALpix sampled (symbols). Right: the product of the number of visits to a HEALpix and its priority, summed over the survey regions of interest and presented as a percentage of the fiducial value expected from an ideal survey.

overall in `baseline_v2.0`, by 8.1%. But this leads to only a small decrease (2.1%) in the summed HEALpix priority, reflecting the strategy's focus on the highest-priority areas.

Of greater concern is the central Bulge diamond region including a limited cross section through the disk and halo populations of the Milky Way, and not providing well-sampled time cadence data for the populations at higher galactic longitudes. These stellar populations, and their dynamics, can differ significantly, and it is particularly important for microlensing to sample different lines of sight through the galaxy, particularly along the arms of the Milky Way. We therefore recommend a less centrally concentrated diamond, plus fields in the Magellanic Clouds, with visits distributed across priority regions at a greater range of longitudes.

We do note significant decreases in the coverage of the Zucker SFR and open clusters in `baseline_v2.0`. This is likely the result of decreased visits to high-longitude Galactic Plane fields, relative to `baseline_v1.5`.

The second GalacticPlaneMetric, "%ofNObsPriority," calculates the product of the number of visits to a given HEALpix and the priority of that pixel, and sums this value over the desired survey regions. We present this metric (Figure 7) as a percentage of the fiducial value that would result from all HEALpix within a given region receiving sufficient visits to detect variables in the category $\tau_{obs} = 5$ days. Though many regions do show an improvement in this metric in `baseline_v2.0`, the Figure of Merit is low for all regions, and decreases overall for the combined science map, despite the increase in visits to the high-priority regions. This reflects the fact that no region is covered at the fiducial cadence in the `baseline_v2.0` strategy.

The goal of these Figures of Merit is to provide a guide to evaluating the trade-offs between cadence and survey area. If no additional time can be dedicated to surveying the Galactic Plane, we recommend redefining the diamond region based on the priority maps described above. We discuss this further in Section 5.

11
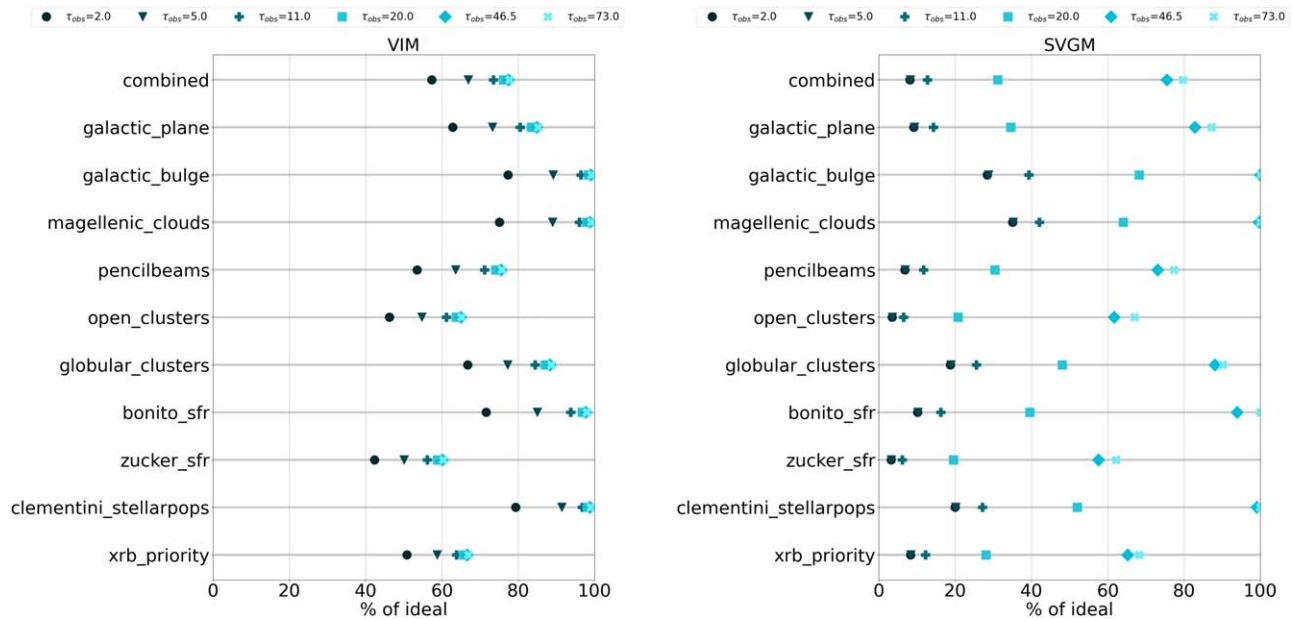
Figure 8. Survey cadence metrics evaluated for the `baseline_v2.0` strategy. Left: survey visit interval metric (VIM). Right: season visit gap metric (SVGM).

It is important to note that one of the Core Community Surveys that will be undertaken by the Nancy Grace Roman Space Telescope will provide very-high-cadence (every ~15 minutes) NIR imaging of a small region (~2 deg^2; Penny et al. 2019) of the central Galactic Bulge. Owing to Roman's viewing and scheduling constraints, these observations will take place in "seasons" ~62–70 days long, interspaced with gaps of several months. Rubin observations are highly complementary to those of Roman in terms of wavelength (optical versus NIR 0.48–2.3 μm), and potentially also in terms of cadence and spatial coverage, if the operational strategies of these contemporaneous surveys are coordinated. For example, Rubin observations during Roman's interseason gaps could help to constrain the parameters of microlensing events (Street et al. 2018a). An in-depth study of Rubin OpSims specifically focusing on microlensing science is in preparation, and a second paper will explore the benefits of combining Roman and Rubin photometry for these events. Rubin's survey of a much wider spatial region is particularly valuable, as it will encompass a greater range of stellar environments and provide information on the populations of those regions that can then be compared with the results of Roman's in-depth study of Bulge populations. This is particularly important for inherently rare categories of variables, such as X-ray binaries (Johnson et al. 2019). We note, however, a few science cases that would not be well served without higher-cadence Rubin observations of at least some Galactic Plane regions, such as CVs and young stellar objects. CVs have blue colors that are not well suited to Roman observations, while young stellar objects exhibit variability on timescales ranging from <10 to 100 days, but are concentrated in SFRs outside the Galactic Bulge.

3.1. Survey Cadence

Figure 8 displays the survey cadence metrics computed for the `baseline_v2.0` strategy for all science regions. In both plots, a value of 100% indicates the ideal cadence over the desired survey region. It is noteworthy that the VIM metric indicates that the `baseline_v2.0` strategy provides a good

cadence for some of the science regions of interest, particularly the Bulge, Magellanic Clouds, and a number of key stellar clusters and SFRs that were not well sampled in earlier strategies. The lack of monitoring of the regions at higher galactic longitude is reflected in the lower value of the metric for all timescales for the combined pencil beams and the X-ray binary science regions.

All but circumpolar fields will experience some inevitable gaps in survey monitoring, due to their annual visibility cycle, and variables (especially transient events) with timescales shorter than this gap will be missed. The SVGM metric takes the inevitable season gaps into account, and instead evaluates whether the gap is longer than expected (the average expected gap between seasons is 145 days). This is primarily of interest for long-timescale phenomena, such as black hole microlensing, for which the lightcurves can be partially sampled if the gap between seasons is elongated. For this reason, we consider this metric for the longer-timescale categories only.

The `baseline_v2.0` strategy maintains relatively uniform sampling across all 10 yr of LSST, and thus Figure 8 indicates that long-timescale variables would be characterized in the Plane, Bulge, Magellanic Clouds, and in many clusters and SFRs. The SVGM metric is mainly intended to compare the impacts of rolling cadence implementations, where large latitudinal bands of the sky may be observed intensively in any given year, but then neglected in favor of a different band in the subsequent year(s), potentially leaving long-term variables uncharacterized. This is discussed further in Section 4.3.

3.2. Filter Selection

Time-dependent color changes are a key diagnostic feature for many variable types. To evaluate whether adequate observations were obtained in different passbands, we calculated the percentage of the desired survey region to receive an R_{exp} ratio of $\geqslant 1.0$. Figure 9 presents this metric for the different survey regions and compares the performance of `baseline_v2.0` and `baseline_v2.1` in different filters. Since both of these strategies apply a uniform exposure time
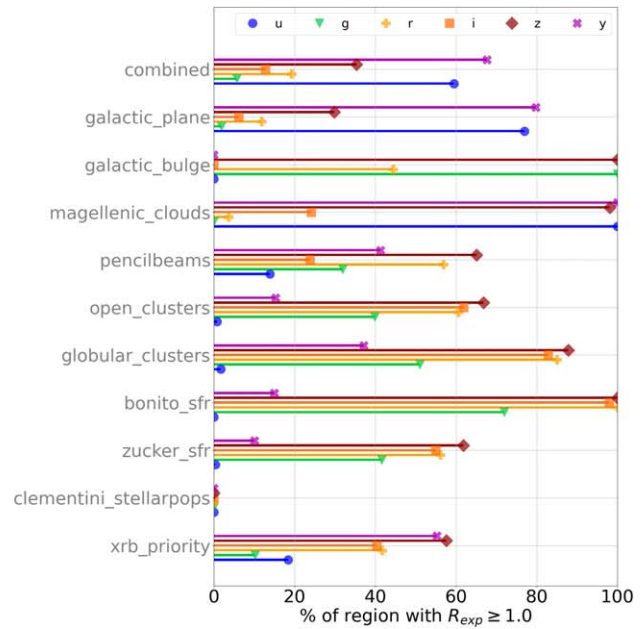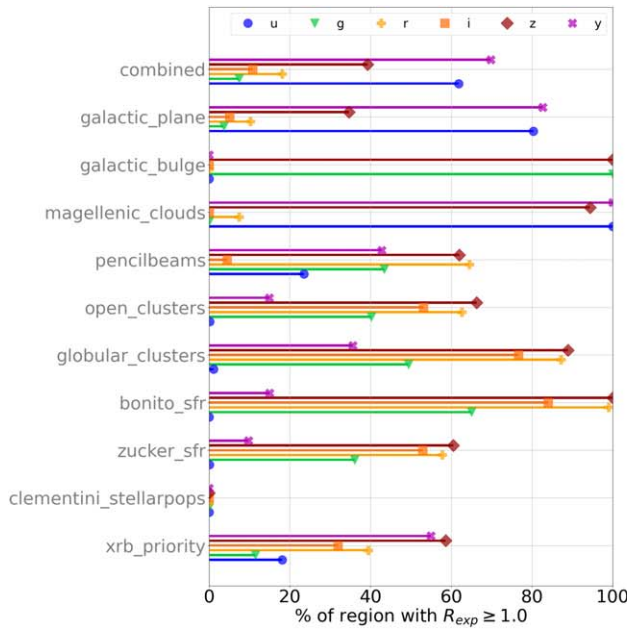
Figure 9. Percentages of desired survey regions to have $R_{exp} \geqslant 1.0$, comparing the time spent dedicated in different passbands for the `baseline_v2.0` (left) and `baseline_v2.1` (right) strategies.

(30 s) in all filters, exposure time spent relates directly to the number of exposures acquired in each filter.

Our analysis indicates that the balance of time spent between different passbands is less than ideal for the combined survey region, for both baseline strategies. This reflects higher sampling in the u and y filters than is recommended for some regions, particularly those with high extinction. Of particular concern is the imbalance between observations in the g band and those in at least one of r, i, or z, since regular observations in both blue and red bands are required for time-series color observations that will identify and characterize many categories of variable stars.

Similar concerns are seen for many of the individual science regions, with the Clementini resolved stellar populations being particularly poorly sampled.

4. Evaluation of v2.0 and v2.1 Simulations

Jones (2021) presented a very large set of simulations in "families" designed to explore how varying different parameters and aspects of the survey strategy would impact different science cases. We evaluated all of the simulations using the metrics described above, and note that this work complements detailed analysis performed for a range of science-specific metrics published in parallel, including the `MicrolensingMetric` (N. S. Abrams et al. 2023, in preparation) and the `NYoungStarsMetric` (Prisinzano et al. 2022a). In particular, the recovery of periodic variables is explored in R. Bonito et al. (2023, in preparation). In this section, we summarize our conclusions. As noted above, the shortest-timescale category of variability is not well served in any simulation, so we present results for $\tau_{var} \geqslant 11$ days.

4.1. vary_gp Family of Simulations

This family of simulations varied the amount of survey time spent on fields in the region of the Galactic Plane not included

in either the Bulge diamond or the (baseline_v2.0) WFD region. Figure 10 summarizes the metric data for this family.

The "%ofPriority" `GalPlaneFootprintMetric` clearly shows that a greater percentage of the priority area is included when more visits are dedicated to the high-galactic-longitude regions of the Galactic Plane, particularly for gpfrac weightings >0.5, with strong increases in the metric for the combined galactic science region and pencil beams. Increasing the sampling of regions outside the central Plane increases the total population of variables of many types, including microlensing events, young stellar objects, and X-ray binaries, which boosts the corresponding metric values when summed over the survey region for longer-timescale ($\tau_{var} > 55$ days) categories of variability.

There is a sweet spot to be identified in terms of cadence and survey footprint. Observing a large area with too low a cadence will mean that transient events cannot be properly identified in time for characterization or follow-up. But reducing the area surveyed in order to boost the cadence will eventually also reduce the total number of transient events discovered. We recommend evaluating the `MicrolensingMetric` with two timescales (30 and 200 days) in order to identify the optimum balance for transient events (see N. S. Abrams et al. 2023, in preparation).

4.2. plane_priority Family of Simulations

This family used the galactic science priority maps described in Section 2.1 to select different regions within the Galactic Plane for inclusion in the WFD, effectively exploring the effects of increasing the size of the diamond region. The "pbf" and "pbt" variants within this family explored the impact of leaving out and including the pencil-beam fields, respectively (see Figure 10, bottom panels).

Unsurprisingly, the footprint metrics (Figure 10) show a marked improvement for the combined survey region when the threshold for HEALpix selection is lowered and a larger region is included; this is also particularly important for coverage of SFRs
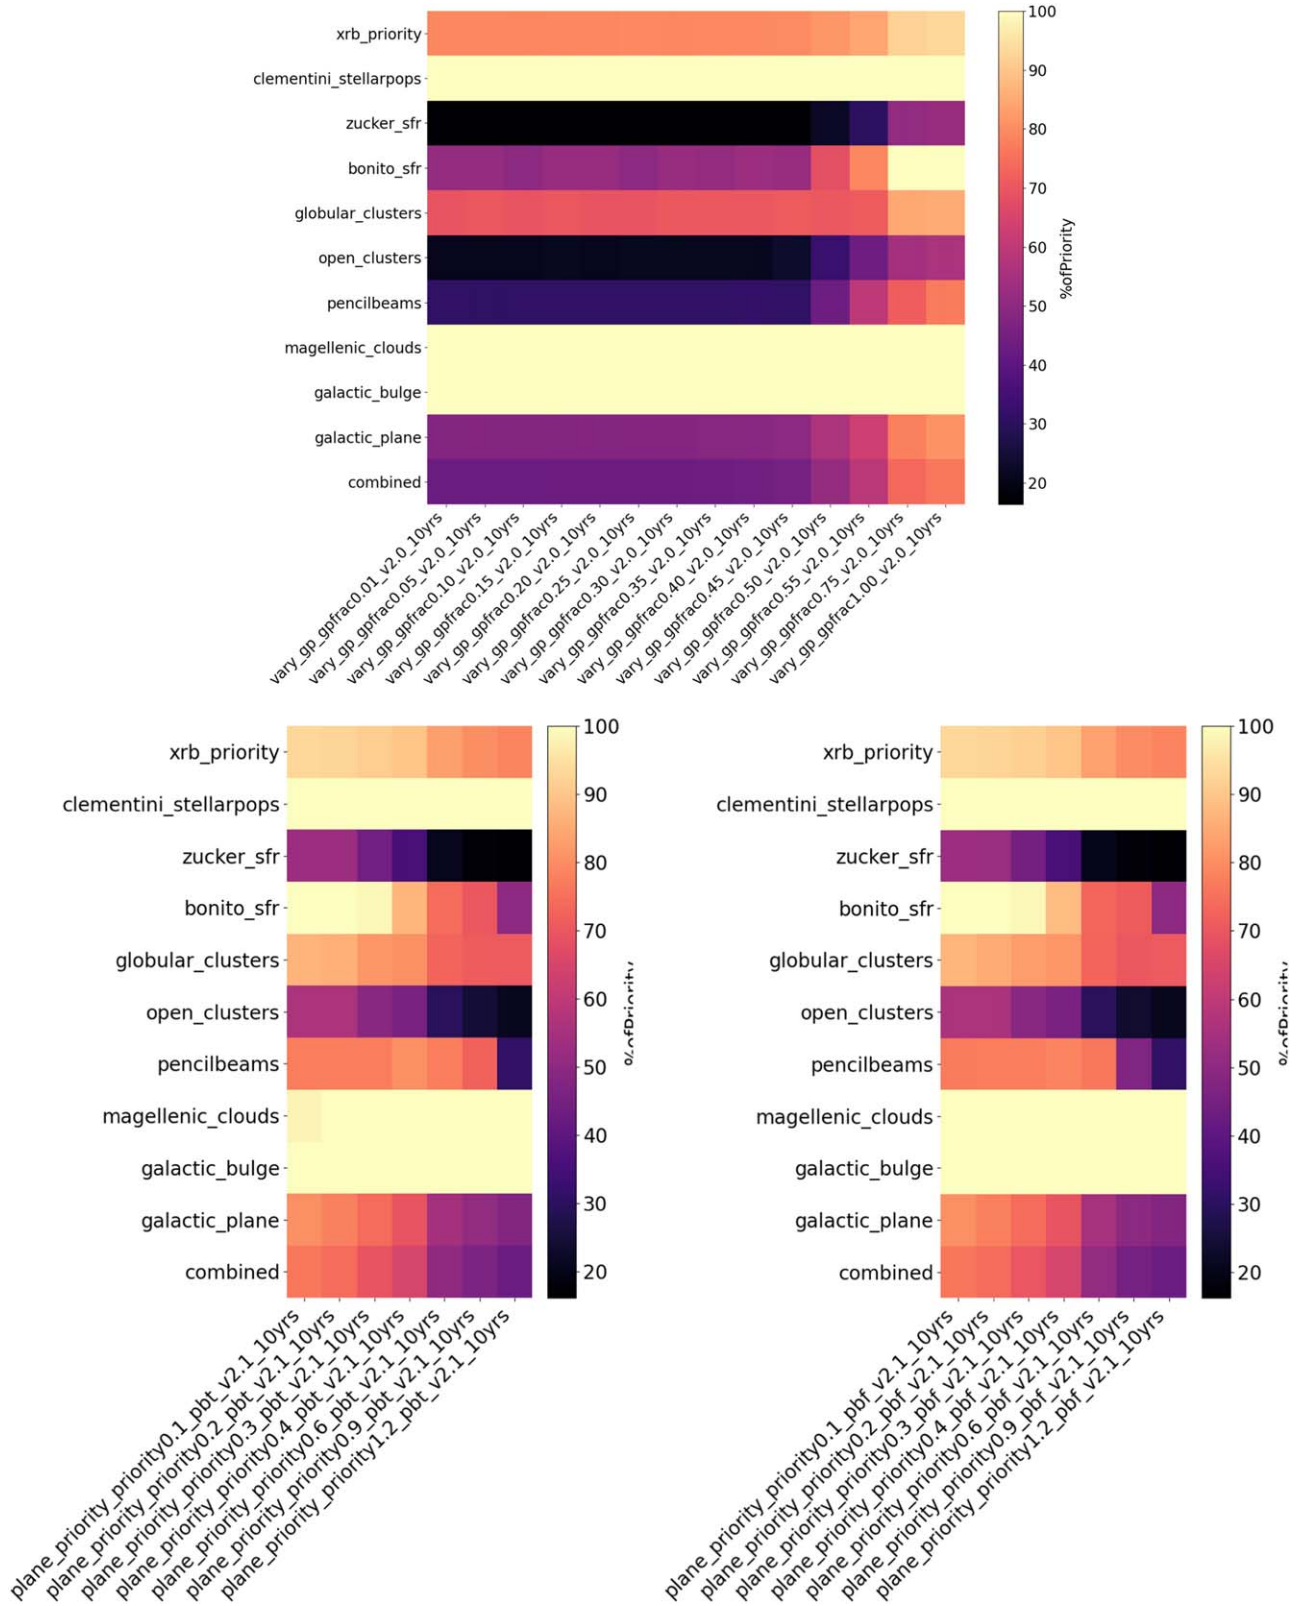
Figure 10. Heatmaps of the "%ofPriority" footprint area monitored at sufficient cadence to characterize variability $\tau_{var} = 11$ days, for simulations in the `vary_gp` (top) and `plane_priority` (bottom) families. In the bottom panels, the left plot summarizes the "pbf" simulations without the pencil-beam fields, while the right plot shows the "pbt" simulations that include them.

and clusters. There is a notable distinction at priority level = 0.9 between the maps with and without the pencil beams included separately. If a priority selection level of 0.6 or lower is used, then the majority of the pencil-beam regions are sampled at a cadence consistent with the rest of the survey region—regardless of whether they are distinctly included in the map or not.
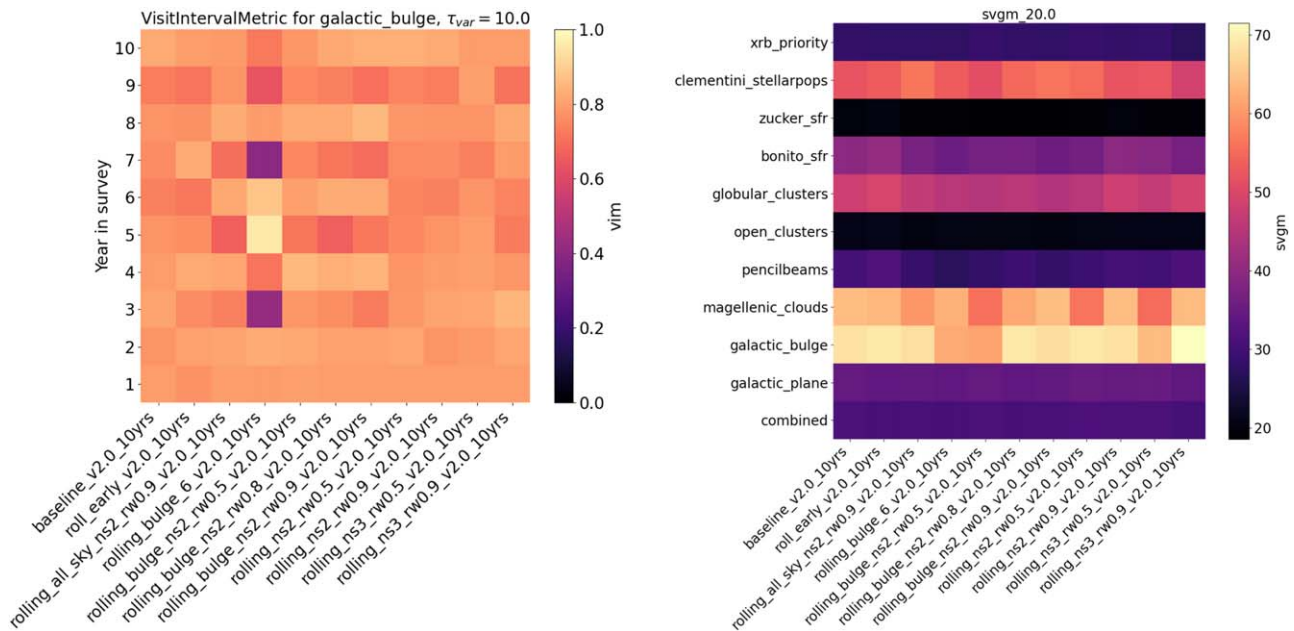
Figure 11. Left: heatmap of the `VisitIntervalMetric` calculated for the Galactic Bulge region for variability $\tau_{var} = 10$ days, for a range of rolling cadence strategies. Right: heatmap of the SVGM calculated for all regions of interest for variability timescales $\tau_{var} = 100$ days.

We found that adopting a priority threshold of at least 0.4 can be used to select a survey footprint that includes the highest-priority regions, while maximizing the cadence achieved.

In terms of observing known Milky Way star clusters to be used as calibrators for the LSST stellar populations, the changes between the different baselines are minimal: `baseline_v2.0` and `baseline_v2.1` observe 70 fewer open clusters in the Galactic Plane, and essentially the same numbers of globular clusters. Nevertheless, there are plenty of open and globular clusters left in all baseline plans ($\sim$2000 and 145, respectively), with a good coverage of the age and metallicity space. There are small improvements in the maximum photometric depth reachable in the latest baselines (up to 0.3 mag in the r band), mostly for clusters in less crowded areas of the Galactic Plane.

4.3. Rolling Cadences

Jones (2021) explored a number of different implementations of a rolling cadence strategy, where the sky is divided into separate bands, with some bands being observed more frequently in alternate years, with the "active" bands switching between regions from year to year. It should be noted that the `baseline_v2.0` includes a two-band rolling cadence for the low-dust WFD region.

Alternatives explored included the use of two, three, or six bands, each with weightings of 0.5, 0.8, or 0.9, referring to the increasing ratio of visits to the active versus inactive bands. Some realizations also applied the rolling cadence to the central Bulge diamond and to the whole sky, while other variations varied in which years the rolling cadence was applied.

Since a similar survey footprint was used in all simulations in this family, we analyzed its impact on the cadence achieved for different timescale categories of variability, in the different galactic science regions of interest. In many cases, no impact was seen, because the rolling cadence was not applied to the Galactic Plane. The exception was the rolling_bulge family of

OpSims, and its impact on the detection of short-timescale variables ($\tau_{var} = 10$ days) can be seen in Figure 11.

Simulation `rolling_bulge_6` (a six-band strategy) indicates that this strategy can reach almost the ideal cadence for the shortest-timescale (2 days) variability category for 1 or 2 yr during the survey, but at the expense of achieving only $\sim$40%–60% of the desired cadence in the remaining years.

If the years of high cadence in the Galactic Bulge were to coincide with the Roman mission's survey of that region, then the complementary data sets would be very valuable for a wide range of science, particularly microlensing planet characterization (Street et al. 2018a). This is a high-priority science case, but it would come at the expense of lower cadences in the other years, which may be detrimental to Rubin's ability to fill in gaps between the Roman observing seasons. Regular observations in all years are necessary to ensure that transient events continue to be discovered over a wide area, even if the cadence is so low that characterization depends on additional follow-up. This also ensures that long-timescale phenomena are properly characterized. We recommend that further exploration of the pros and cons of rolling cadence in the Bulge be undertaken. This should include comparing simulations of rolling cadence being implemented for just the very small Roman Bulge survey region with those of rolling cadence over a larger area, while a regular cadence is maintained elsewhere in the Galactic Plane. At $\sim$2 deg^2, this survey region spans just $\sim$20% of a single Rubin field of view, so a Rubin deep-drilling field (DDF) at this location could not only deliver complementary wavelength and cadence photometry, but would also provide data on the wider surrounding region, helping to put the Roman results in context.

The simulation of a rolling cadence across the whole sky is also noteworthy. In high-cadence years, it delivers $\sim$63% of the desired sampling for the shortest-timescale variable categories in the Galactic Plane region. In "off" years, this drops to $\sim$60%, but the improvement in sampling over the baseline ($\sim$62%) was small. We used the SVGM to explore whether the rolling cadence strategies created any undesirably long gaps between

annual seasons of observations of the different regions of interest (Figure 11). For variability timescales $\tau_{var} \geqslant 100$ days, the metric changed by $<10\%$, suggesting that no gaps in the lightcurves would be detrimental to the characterization of long-timescale variables.

4.4. Intranight Cadences

A number of LSST White Papers recommended multiple visits per night in various regions, to improve the rejection of spurious transient candidates and the identification of moving objects. The `presto` and `long_gaps` simulations explored strategies where two or three visits per night were performed, with different time intervals between them, as well as repeated or different filter selections.

However, most of these simulations applied these modifications to the WFD survey region. There was a significant ($\sim 30\%$) reduction in the area covered at adequate cadence for $\tau_{var} \geqslant 55$ days in the Magellanic Clouds if the gap between exposures is short, i.e., <3 hr. There was a similar reduction in cadence across open and globular clusters, though the impact was less marked. Otherwise, the results were very similar to the `baseline_v2.0` strategy, and shorter timescales of variability are not well sampled.

Evaluating the balance of time allocated to different filters showed a dramatic variation in the results for different regions of science interest in different filters. For example, sufficient data were obtained in g band for the Galactic Bulge in all simulations, but the results in this filter for the Magellanic Clouds, Galactic Plane, X-ray binaries, and resolved stellar populations were very poor. Conversely, the results in z band were much better, with most regions receiving at least 60% of the desired data.

In r band, the "non-mixed" filter pairings among the presto_gap strategies generally returned the best combination of filter and cadence, since these focused on the g, r, i, z filter set recommended for most galactic science. The achievable cadence in r for the "mixed" filter set g, r, i, z, y was considerably lower in all science regions ($<20\%$ of desired). However, this was reversed in the i band, where an improved cadence was found with the mixed filter set in most presto_gaps implementations.

We note that pairs of exposures in different filters are very useful for distinguishing variability classes in real time, and this is particularly important for brokers to be able to identify transients, such as microlensing and CV outbursts. The characterization of microlensing will depend on our ability to constrain the source magnification as a function of time in at least two colors during the event. Regular multicolor observations are also important for the characterization of pulsating stars and other periodic variables.

We recommend that different filter sets be prioritized with independent cadences separately for the WFD and galactic science regions of interest, respectively.

4.5. Other Simulations Explored

Jones (2021) included a wide range of other simulations designed to explore different scenarios not directly relevant to galactic science. Nevertheless, it was important to assess these simulations as well, to identify any unexpected negative impacts. Here we briefly summarize our key findings.

1. `Bluer Balance`: Increasing the number of g-band exposures is reflected by an increase in the VIM for SFRs. Enhancing both u and g exposures causes a detrimental decrease in cadence for almost all regions.
2. `Long u`: Increasing u-band exposures to 50 s exposure time enhanced the percentage of some regions of interest that received the desired ratio of visits in u band, provided that the total number of u-band exposures in the survey remained the same. X-ray binaries, the Galactic Bulge, and pencil-beam fields benefited. However, this came at the expense of the number of observations in the other filters that are important for other science goals.
3. `Vary Exposure Time/Shave`: This strategy slightly reduced the area of the desired footprint covered at a cadence sufficient to detect variability timescales of 55 days or longer, negatively impacting primarily science from the Galactic Plane, open clusters, and SFRs. However, the VIM per year shows significant improvement—1.0 (max) versus 0.8—over `baseline_v2.0` for longer-variability ($\tau_{var} > 55$ days) timescales. This may be due to a higher number of observations reaching the required signal-to-noise thresholds, thereby improving the cadence achieved.
4. `Vary North Ecliptic Spur`: By including more visits to fields at high galactic longitude, more SFRs are covered at higher cadence. The simulations with a weighting factor nesfrac $\geqslant 0.75$ show a significant improvement for this science case. However, it appears to come at the expense of coverage in the Galactic Plane, reducing the cadence in the combined region of interest by $\sim 10\%$.
5. `No Repeat`: The imposition of an additional basis function to avoid repeated visits in a single night seems to slightly reduce the number of visits to the Galactic Plane relative to `baseline_v2.1`, evidenced by a drop in the area receiving the cadence required in each time category. This is detrimental to many of our science goals. However, simulations with a higher repeat weight parameter do show a strong improvement in the SeasonGaps metric, specifically for the Galactic Bulge region, which is beneficial for long-term variable classes in that region.
6. `Good Seeing`: Requiring that g, r, i images be acquired in good seeing conditions for the whole sky results in a small reduction in the areas of interest receiving adequate cadence, relative to `baseline_v2.1`. Given the importance of these images to good image subtraction, particularly in crowded regions like the Galactic Plane and Magellanic Clouds, we conclude that the benefits outweigh the detriments.
7. `Deep-drilling Fields`: Reducing the number of visits to the DDFs to 3% causes a small but unsurprising increase in the cadence in galactic science regions of interest, with a corresponding decrease if more time is dedicated to the DDFs.

Since none of the DDFs lie within the galactic science regions of interest, none of the DDF simulations have any impact on our footprint or cadence metrics. Some of the DDF "accordion" simulations did produce undesirable variations in the frequency of r-band monitoring in the Galactic Bulge specifically, underscoring the importance of implementing different cadences per filter for different survey regions.

4.6. Microsurveys

This category of simulations includes a wide range of proposals made in general to support specific science cases, often with specialized constraints on the survey region or cadence. To some degree, whether adequate on-sky time will be available for these proposals to be included in LSST after the science requirements are met depends on the hardware overheads relating to slew times, filter changes, etc. Current estimates suggest that some time will be available, but a final decision on the microsurveys will be made after commissioning. That said, a number of the microsurvey proposals are particularly relevant to galactic science, so we reviewed the simulations provided and summarize our findings below.

1. Virgo Cluster: This region lies outside the desired regions for galactic science, and its inclusion had no impact on the cadence achieved within our desired regions.

2. SMC movie: The requested two nights of intensive observations of this relatively small region have no appreciable impact on the metrics for the other regions, while the survey would enable the detection of exoplanets in the SMC and characterize short-term variables of all kinds. The alternative strategy proposed (longer-cadence observations over a longer period) would be delivered automatically by including the SMC in the WFD region, as done for baseline_v2.0 and _v2.1.

3. Roman: This strategy proposed a dual cadence: high-cadence observations of the Bulge within limited periods where Roman simultaneously observes the same Galactic Bulge Time Domain Survey field, and lower-cadence observations between the Roman survey windows. Some of the lower-cadence observations would be provided automatically if the Bulge field is included in the WFD region, as in baseline_v2.0, _v2.1. VIM shows a small increase for the Bulge region, reflecting the relatively small number of additional observations required for the high-cadence strategy, with no significant impact on other regions.

4. Too_rate: Rubin performing 10–50 target-of-opportunity overrides per year appears to have no significant impact on the observations in our regions of interest.

5. North Stripe: This strategy would increase the number of visits to regions of the Galactic Plane at high longitude, which is beneficial for some galactic science.

6. Short exposure: This strategy has no appreciable impact on our metrics, meaning that the additional exposures do not subtract a significant number of visits per region overall. We note the value of these short exposures for calibrating Rubin's galactic star map at the bright end of its magnitude range, which will allow the stellar populations to be more easily compared between Rubin data and other surveys (Gizis 2018).

7. Multi_short: Taking four short exposures back to back in a single filter per pointing (up to a maximum of $12\,\mathrm{yr}^{-1}$) slightly reduced the cadence achieved in almost all of our desired survey regions, except the Galactic Bulge.

8. Twilight NEO: These simulations prioritized observations in twilight of areas outside the Galactic Plane, where near-Earth objects are easiest to detect. This produces a slightly lower cadence (relative to baseline_v2.0) in all years for all regions of interest, notably the pencil beams, SFRs, and the Galactic Plane.

9. Carina: These simulations were designed around the Bonito & Hartigan (2018) proposal. The v1.7 simulation represented the best strategy of one visit to the Carina Nebula every 30 minutes, cycling through filters g, r, i, and u for seven days (with this observation pattern being repeated every 2–3 yr). The goal of the proposal was to execute this pattern for five SFRs, including Carina. Our metrics indicate this strategy significantly enhanced the VIM for the Bonito SFRs, as expected (to 95% for $\tau_{\mathrm{var}} = 11$ days compared with 93% for the baseline_v2.0), with no impact to other regions. A more specialized analysis of the characterization of accreting young stellar objects is presented in Bonito & Venuti (2021).

5. Polishing the Diamond: Revised Survey Region in the Galactic Plane

Although the inclusion of the diamond in the central Galactic Plane and the Magellanic Clouds in the baseline_v2.0 and _v2.1 proved to be a dramatic improvement for many of the science cases considered here, the boundaries of the diamond were quite approximate in the first round of simulations. In this section, we describe how the priority maps presented in Section 2.1 were used to refine the diamond region to better reflect areas of high scientific interest.

First, rather than include all galactic clusters and SFRs in the footprint at the same priority, we implemented a relative priority weighting for the HEALpix within different regions, based on the recommendations of specialists in those fields of study. Open clusters were ranked in priority according to the number of Gaia stars identified within the cluster (L. Giradi 2023, private communication). SFRs were prioritized according to the order of Table 3 in Prisinzano et al. (2022b).

The maps can be used to directly select larger or smaller regions, simply by selecting those HEALpix with a priority above a lower or higher threshold, respectively. However, the resulting footprint tends to break into numerous smaller regions that are often separated by gaps larger than the Rubin field of view, unless a very low threshold is used to select a very large area. This is a natural consequence of the localized nature of some areas of interest (e.g., star clusters) and the nonhomogeneous extinction due to patchy dust across the Plane.

The patchy footprint raised the concern that this would cause the LSST scheduler to spend a higher amount of time slewing between regions rather than exposing, decreasing the survey efficiency and the total number of exposures possible over the survey as a whole.

To explore this possibility, two simulations were compared that implement alternative sets of pencil-beam fields to cover regions of high interest in the Galactic Plane (Table 3). The pencil_fs1 simulation added 20 single-pointing pencil beams, distributed across the Galactic Plane, to the baseline_v2.0 footprint, whereas pencil_fs2 added an alternative set of four larger pencil beams. Both sets of pencil beams include the same number of square degrees of sky in total.

Two metrics were used to evaluate survey efficiency: the number of visits realized to the science regions of interest and the OpenShutterFractionMetric. We found that larger,
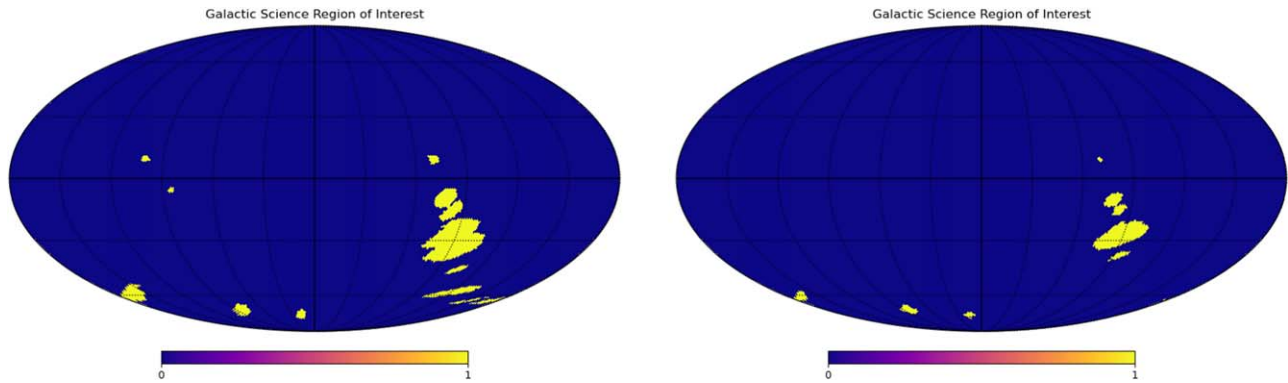
Figure 12. Maps of the proposed revised survey regions in the Galactic Plane and Magellanic Clouds in Mollweide projections. Left: wider region selected with a projected pixel priority threshold of 1.5. Right: smaller region selected with a threshold of 2.0.

consolidated pencil beams result in a 7% increase to these regions overall, with little change to the open shutter fraction in pencil-beam fields. Most of the other regions of interest show small (∼2%) increases or decreases in their number of visits, depending on whether they overlapped the pencil beams. However, coverage of the Bonito SFRs dropped considerably —receiving 39% fewer visits—when larger pencil beams were used.

In general, the changes in the open shutter fraction were small for all regions, but the standard deviation of the shutter fraction within each science region of interest changed by an order of magnitude for some regions. Within the pencil-beam regions, the stddev dropped from 0.16 to 0.02 between the smaller and larger sets, respectively. This suggests that the overheads to image different HEALpix within those regions were more consistent.

We concluded that survey efficiency can be improved by selecting larger, contiguous regions for galactic science wherever possible. In order to design this region to maximize the science as fairly as possible across the different science cases, we adopted the following approach.

The original priority maps (Figure 2) display numerous spatial regions of HEALpix ("peaks") at high priority when plotted in a Mollweide sky projection. The Mollweide coordinates of the centroids of these regions were identified using the photutils.find_peaks function. We consider each peak as a potential pointing for the Rubin Observatory, and identify those pixels that would lie within a single Rubin field of view. Here a selection threshold was applied, to include only pixels with a priority above a given threshold; this allows us to select larger or smaller survey regions by setting the priority threshold. We refer to the resulting pixel set as a "cluster."

The goal was to identify "superclusters," where the pixel clusters are sufficiently close together, relative to the angular scale of the Rubin field of view, that they can be merged to form a region encompassing several high-priority areas. The superclusters including the greatest cumulative priority value (summed over their HEALpix area) can then be added to a more contiguous, revised footprint region.

Numerous algorithms exist to identify clusters in parameter space, so we first explored a range of algorithms implemented by the sklearn.cluster package, including DBSCAN (Ester et al. 1996), MiniBatchKMeans (Arthur & Vassilvitskii 2007), OPTICS (Ankerst et al. 1999), MeanShift (Comaniciu & Meer 2002), and AffinityPropagation (Frey & Dueck 2007). We experimented with feature sets, including a

matrix of the separations between the peaks and their priority values, and tuned the configurable parameters of each algorithm to optimize the results. Of these, we found that AffinityPropagation and MiniBatchKMeans produced the "best" results. This was judged by comparing the revised survey footprint produced, with the criteria that it should consist of a small number of contiguous regions, which needed to include the highest-priority peaks and clusters visible on the maps, without also including large areas of low-priority pixels.

However, this approach emphasized spatial proximity features over the HEALpix priority, which led to a number of smaller (but high-priority) regions of the sky being neglected, such as the Small Magellanic Cloud. We therefore adopted an alternative approach. We identified potential superclusters where multiple clusters lay within a region four times the angular width of the Rubin field of view. The boundaries of each candidate supercluster were defined as ellipses centered on the mean (x, y) of the member clusters, with semimajor axes set by the minimum and maximum x, y ranges of the included clusters. For a candidate supercluster to be accepted, the sum of the priorities of all pixels within the ellipse was required to be greater than the sum of the pixels in the member clusters. This test prevented superclusters from expanding continuously into low-priority regions.

This approach produced a more consolidated survey footprint, while the use of elliptical superclusters minimized the number of low-priority areas included. It also allowed us to reweight the relative priority given to each science case, to ensure that science that focuses on specific small regions was not dominated by science requiring larger areas.

The original diamond included a region of ∼400 deg². To identify a more optimized survey region, we first used the above process to simply redistribute this same total area over high-priority regions. We computed the total area selected by the above process for priority thresholds from 1.0 to 4.4, at intervals of 0.1, and found that a threshold of 2.0 selected a region of 424.7 deg².

However, we noted that this region excludes a number of regions of high priority for several science cases, including several pencil beams and SFRs, so we next identified the smallest survey region that included these areas. We found that a selection threshold of 1.5 selected an area totaling 937.5 deg² that included this minimum survey region. Both of the resulting survey regions were proposed for the next round of survey strategy simulations, and are shown in Figure 12. We stress that these regions represent two possible consolidated survey

footprints, but alternative regions can be achieved by adjusting the selection thresholds used. For example, we recommend exploring the potential benefits of applying a rolling cadence to very limited regions, including the Roman Galactic Bulge Time Domain Survey field, Magellanic Clouds, and selected SFRs, where high-cadence data are most valuable. By restricting this high-cadence region, it may be possible to maintain a regular cadence over a large area of the rest of the Galactic Plane, thereby balancing the needs of different science cases.

6. Conclusions

We analyzed all of the `_v2.0` and `_v2.1` Rubin survey strategy simulations to evaluate how well they performed for time-domain galactic science, with particular reference to transient phenomena. While other works have explored the simulations with specific science goals, the objective of this work was to provide high-level metrics to assist with optimizing Rubin's survey strategy for as broad a range of science as possible. We approached this by considering the three major observational factors in survey strategy design— sky region, cadence, and the balance of exposures between filters. We combined the spatial regions of interest for a range of different galactic science to produce a map indicating the relative priority per HEALpix, and used this to recommend the most scientifically valuable regions of the Galactic Bulge, Plane, and Magellanic Clouds to include in LSST. A wide range of time-domain galactic science can be represented by categorizing variability according to timescale, but this needs to be considered together with survey footprint, since the spatial distributions of the target objects are often highly localized (e.g., star clusters).

6.1. Key Findings and Recommendations

`baseline_v2.0` and `_2.1` were shown to deliver significantly improved cadences, relative to `_v1.5`, in all of the science cases considered, thanks to the inclusion of a region in the central Galactic Plane and the Magellanic Clouds in the WFD. However, in our analysis of `_v2.0` and `_v2.1` simulations, we concluded that very little of the regions of interest to galactic science would receive sufficient visits to reliably detect short-timescale (<10 days) variability, unless a rolling cadence strategy is adopted for the Galactic Plane (over the whole or a subregion). In particular, we note that characterizing short-term variability in young stellar objects will require a higher cadence in specific regions outside the Galactic Bulge, as discussed in Bonito & Venuti (2021). We propose a "refined diamond," consisting of a limited footprint of interest where high-cadence observations would be most scientifically valuable, but emphasize that the rest of the Galactic Plane should still receive regular visits throughout the survey. This ensures that long-term variability can be characterized across the whole Galactic Plane, and that the coadded images will reach sufficiently faint limiting magnitudes to ensure a large-volume survey for SFRs and stellar clusters. We recommend further simulations of a rolling cadence strategy applied to the refined diamond regions, particularly to coordinate the rolling seasons with the Roman Galactic Bulge Time Domain Survey, to enhance the characterization of microlensing events from that region. Last, we note that, for galactic science, there are pros and cons to the choice between 1×30 s exposures (as adopted in `baseline_v1.5`) and 2×15 s for the g, r, i, z, y filters (used in `baseline_v2.0`). Taking fewer, longer exposures results in lower operational overheads, and theoretically allows for $\sim 9\%$ more visits overall, at least some of which could be used to either enhance the cadence or area covered in the Galactic Plane. On the other hand, most of the science cases considered here depend on the cadence of observations, rather than reaching deep limiting magnitudes, and shorter exposures provide more unsaturated photometry of brighter objects. This allows for the comparison of LSST photometry with that of other surveys, as well as facilitating follow-up observations.

ORCID iDs

R. A. Street https://orcid.org/0000-0001-6279-0552
X. Li https://orcid.org/0000-0002-0514-5650
S. Khakpash https://orcid.org/0000-0002-1910-7065
E. Bellm https://orcid.org/0000-0001-8018-5348
L. Girardi https://orcid.org/0000-0002-6301-3269
L. Jones https://orcid.org/0000-0001-5916-0031
N. S. Abrams https://orcid.org/0000-0002-0287-3783
Y. Tsapras https://orcid.org/0000-0001-8411-351X
E. Bachelet https://orcid.org/0000-0002-6578-5078
P. Gandhi https://orcid.org/0000-0003-3105-2615
P. Szkody https://orcid.org/0000-0003-4373-7777
W. I. Clarkson https://orcid.org/0000-0002-2577-8885
R. Szabó https://orcid.org/0000-0002-3258-1909
L. Prisinzano https://orcid.org/0000-0002-8893-2210
R. Bonito https://orcid.org/0000-0001-9297-7748
R. Di Stefano https://orcid.org/0000-0003-0972-1376

References

Ankerst, M., Breunig, M. M., Kriegel, H.-P., & Sander, J. 1999, in Record of the ACM Special Interest Group on the Management of Data (SIGMOD), ed. S. Davidson & C. Faloutsos (New York: Association for Computing Machinery), 49
Arthur, D., & Vassilvitskii, S. 2007, in SODA '07: Proc. Eighteenth Annual ACM-SIAM Symp. on Discrete Algorithms, ed. H. Gabow (Philadelphia, PA: SIAM), 1027
Baumgardt, H., & Hilker, M. 2018, MNRAS, 478, 1520
Bianco, F. B., Ivezić, Ž., Jones, R. L., et al. 2022, ApJS, 258, 1
Bonito, R., & Hartigan, P. 2018, Young Stars and Their Variability with LSST, LSST White Paper 2018, https://docushare.lsstcorp.org/docushare/dsweb/Get/Document-30505/bonito_carina_dd.pdf

Bonito, R., & Venuti, L. 2021, Young Stellar Objects and Their Variability with Rubin Observatory LSST, LSST Survey Cadence Note, https://docushare.lsst.org/docushare/dsweb/Get/Document-37625/rbonito.pdf

Bonito, R., Venuti, L., Ustamujic, S., et al. 2023, ApJS, 265, 27

Bono, G., Dall'Ora, M., Fabrizio, M., et al. 2018, Unveil the Darkness of the Galactic Bulge (VESTALE), LSST White Paper 2018, https://docushare.lsstcorp.org/docushare/dsweb/Get/Document-30571/dalla_ora_vestale_gp.pdf

Clementini, G., Musella, I., Chieffi, A., et al. 2018, The Gaia-LSST Synergy: Resolved Stellar Populations in Selected Local Group Stellar Systems, LSST White Paper 2018, https://docushare.lsstcorp.org/docushare/dsweb/Get/Document-30585/clementini_stellarpop_wfd.pdf

Comaniciu, D., & Meer, P. 2002, ITPAM, 24, 603

Dal Tio, P., Pastorelli, G., Mazzi, A., et al. 2022, ApJS, 262, 22

Dehnen, W., & Binney, J. 1998, MNRAS, 294, 429

Di Criscienzo, M., Leccia, S., Braga, V., et al. 2023, ApJS, 265, 41

Ester, M., Kriegel, H.-P., Sander, J., & Xu, X. 1996, in Proc. Second Int. Conf. on Knowledge Discovery and Data Mining, KDD'96, ed. E. Simoudis, J. Han, & U. Fayyad (Washington, DC: AAAI Press), 226

Frey, B. J., & Dueck, D. 2007, Sci, 315, 972

Geller, A., Polsin, A., Bowen, A., & Miller, A. 2021, ApJ, 919, 83

Gizis, J. 2018, Calibrating Milky Way Maps: An LSST Bright(ish) Star Survey, LSST White Paper 2018, https://docushare.lsstcorp.org/docushare/dsweb/Get/Document-30579/gizis_brightstar_minisurvey.pdf

Gonzalez, O. A., Clarkson, W., Debattista, V. P., et al. 2018, The Definitive Map of the Galactic Bulge, LSST White Paper 2018, https://docushare.lsstcorp.org/docushare/dsweb/Get/Document-30589/gonzalez_stellarpops_gp.pdf

Gorski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759

Grimm, H. J., Gilfanov, M., & Sunyaev, R. 2002, A&A, 391, 923

Ivezic, Z. & LSST Science Collaboration 2018, The LSST System Science Requirements Document LPM-17, https://docushare.lsstcorp.org/docushare/dsweb/Get/LPM-17

Johnson, C. I., Rich, R. M., Young, M. D., et al. 2020, MNRAS, 499, 2357

Johnson, M. A. C., Gandhi, P., Chapman, A. P., et al. 2019, MNRAS, 484, 19

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

Kharchenko, N. V., Piskunov, A. E., Schilbach, E., Röser, S., & Scholz, R. D. 2013, A&A, 558, A53

Lund, M., Stassun, K., Farihi, J., et al. 2018, A Higher Cadence Subsurvey Located in the Galactic Plane, LSST White Paper 2018, https://docushare.lsstcorp.org/docushare/dsweb/Get/Document-30602/lund_transits_gp.pdf

Moniez, M., Sajadian, S., Karami, M., Rahvar, S., & Ansari, R. 2017, A&A, 604, A124

Olsen, K., Szkody, P., Cioni, M.-R., et al. 2018, Mapping the Periphery and Variability of the Magellanic Clouds, LSST White Paper 2018, https://docushare.lsstcorp.org/docushare/dsweb/Get/Document-30645/olsen_mc_mini.pdf

Penny, M. T., Gaudi, B. S., Kerins, E., et al. 2019, ApJS, 241, 3

Poleski, R., & Mróz, P. 2018, The First Extragalactic Exoplanets—What We Gain from High Cadence Observations of the Small Magellanic Cloud?, LSST White Paper 2018, https://docushare.lsstcorp.org/docushare/dsweb/Get/Document-30584/poleski_smc_mini.pdf

Prisinzano, L., Carnerero, M. I., Balmaverde, B., et al. 2022a, ApJS, 258, 3

Prisinzano, L., Damiani, F., Sciortino, S., et al. 2022b, A&A, 664, A175

Prisinzano, L., & Magrini, L. 2018, Investigating the Population of Galactic Star Formation Regions and Star Clusters within a Wide-Fast-Deep Coverage of the Galactic Plane, LSST White Paper 2018, https://docushare.lsstcorp.org/docushare/dsweb/Get/Document-30483/prisinzano_starformation_gp.pdf

Raiteri, C. M., Carnerero, M. I., Balmaverde, B., et al. 2022, ApJS, 258, 3

Strader, J., Aydi, E., Britt, C., et al. 2018, The Plane's the Thing: The Case for Wide-Fast-Deep Coverage of the Galactic Plane and Bulge, LSST White Paper 2018, https://docushare.lsstcorp.org/docushare/dsweb/Get/Document-30482/strader_plane_wfd.pdf

Street, R., Lund, M., Donachie, M., et al. 2018a, arXiv:1812.04445

Street, R., Lund, M., Khakpash, S., et al. 2018b, arXiv:1812.03137

Suberlak, K., Slater, C., & Ivezić, b. 2018, LSST Fall 2017 Crowded Fields Testing DMTN-077, https://dmtn-077.lsst.io

Tan, X., & Showman, A. P. 2019, ApJ, 874, 111

Tsapras, Y., Street, R. A., Hundertmark, M., et al. 2019, PASP, 131, 124401

Wiktorowicz, G., Middleton, M., Khan, N., et al. 2021, MNRAS, 507, 374

Zucker, C., Speagle, J. S., Schlafly, E. F., et al. 2020, A&A, 633, A51