



Publication Year	2021
Acceptance in OA	2022-03-22T13:19:35Z
Title	Blazar Variability with the Vera C. Rubin Legacy Survey of Space and Time
Authors	RAITERI, Claudia Maria, Carnerero, Maria I., BALMAVERDE, Barbara, Bellm, Eric C., Clarkson, William, D'AMMANDO, FILIPPO, Paolillo, Maurizio, Richards, Gordon T., VILLATA, Massimo, Yoachim, Peter, Yoon, Ilsang
Publisher's version (DOI)	10.3847/1538-4365/ac3bb0
Handle	http://hdl.handle.net/20.500.12386/31778
Journal	THE ASTROPHYSICAL JOURNAL SUPPLEMENT SERIES
Volume	258



Blazar Variability with the Vera C. Rubin Legacy Survey of Space and Time

Claudia M. Raiteri¹, Maria I. Carnerero¹, Barbara Balmaverde¹, Eric C. Bellm², William Clarkson³, Filippo D'Ammando⁴, Maurizio Paolillo⁵, Gordon T. Richards⁶, Massimo Villata¹, Peter Yoachim⁷, and Ilsang Yoon⁸

¹ INAF-Osservatorio Astrofisico di Torino, Via Osservatorio 20, I-10025 Pino Torinese, Italy; claudia.raiteri@inaf.it

² DIRAC Institute, Department of Astronomy, University of Washington, 3910 15th Avenue NE, Seattle, WA 98195, USA

³ University of Michigan-Dearborn, Dearborn, MI, USA

⁴ INAF-Istituto di Radioastronomia, Via Gobetti 101, I-40129 Bologna, Italy

⁵ Università degli Studi di Napoli Federico II, I-80126 Napoli, Italy

⁶ Department of Physics, Drexel University, 32 S 32nd Street, Philadelphia, PA 19104, USA

⁷ University of Washington, Seattle, WA, USA

⁸ National Radio Astronomy Observatory, Charlottesville, VA, USA

Received 2021 October 15; revised 2021 November 10; accepted 2021 November 17; published 2021 December 22

Abstract

With their emission mainly coming from a relativistic jet pointing toward us, blazars are fundamental sources for studying extragalactic jets and their central engines, consisting of supermassive black holes fed by accretion disks. They are also candidate sources of high-energy neutrinos and cosmic rays. Because of the jet orientation, the nonthermal blazar emission is Doppler beamed; its variability is unpredictable, and it occurs on timescales from less than 1 hr to years. Comprehension of the diverse mechanisms producing the flux and spectral changes requires well-sampled multiband light curves over long time periods. In particular, outbursts are the best test bench for shedding light on the underlying physics, especially when studied in a multiwavelength context. The Vera C. Rubin Legacy Survey of Space and Time (Rubin-LSST) will monitor the southern sky for 10 yr in six photometric bands, offering a formidable tool for studying blazar variability features in a statistical way. The alert system will allow us to trigger follow-up observations of outstanding events, especially at high (keV-to-GeV) and very high (TeV) energies. We here examine the simulated Rubin-LSST survey strategies with the aim of understanding which cadences are more suitable for blazar variability science. Our metrics include light curve and color sampling. We also investigate the problem of saturation, which will affect the brightest and many flaring sources, and will have a detrimental impact on follow-up observations.

Unified Astronomy Thesaurus concepts: Active galactic nuclei (16); Blazars (164); Flat-spectrum radio quasars (2163); BL Lacertae objects (158); Plasma jets (1263)

1. Introduction

At the centers of many galaxies, supermassive black holes (SMBH) of millions to billions of solar masses are fed by matter falling from surrounding accretion disks, releasing gravitational energy. These active galactic nuclei (AGN) show a variety of different properties, which also depend on their orientation with respect to the line of sight (e.g., Urry & Padovani 1995). In radio-loud objects, two plasma jets are emitted (nearly) perpendicular to the disk. When one of the jets is oriented near to the line of sight, its emission is relativistically Doppler beamed, and it usually dominates over all of the other sources of radiation. The objects where this occurs are called “blazars.” They show random flux variability at all wavelengths, from the radio band to the γ -rays, with timescales ranging from less than 1 hr to years. In general, low-amplitude fast variations overlap with larger and slower flux oscillations, indicating that different variability mechanisms are at work. Spectral variability is usually also detected, sometimes with definite trends versus brightness.

The Vera C. Rubin Legacy Survey of Space and Time (Rubin-LSST) offers a magnificent opportunity to study the whole population of blazars in the Southern Hemisphere and beyond, and to address the still open questions about their variability,

census, and environment. The 8.4 m telescope in Chile will scan more than 18,000 square degrees of sky for 10 yr in the six bands u , g , r , i , z , and y (Ivezić et al. 2019). The field of view will be 9.6 square degrees. At least 80% of the Rubin-LSST time will be dedicated to the main survey, the Wide-Fast-Deep (WFD) survey. Four fields will be observed at a much higher rate: the deep drilling fields (DDFs) will cover regions of the sky that have already been the subject of intensive multiwavelength observing efforts. A further double DDF will match the Euclid Deep Field South.⁹ Moreover, three minisurveys will likely explore the Galactic Plane, the North Ecliptic Spur, and the South Celestial Pole, though with a lower cadence than the WFD.

Through the study of blazar variability, we aim to understand what happens in extragalactic jets at (sub)parsec scales. There, particles are accelerated at relativistic speeds, emitting synchrotron radiation observed in the radio-to-UV¹⁰ frequency bands, and X-ray and γ -ray photons through inverse Compton and/or hadronic processes (e.g., Böttcher et al. 2013). Blazars are formidable cosmic accelerators that are also alleged to produce high-energy neutrinos and cosmic rays (e.g., Murase et al. 2012; Aartsen et al. 2018a, 2018b; Giommi et al. 2020). Moreover, through the analysis of blazar variability, we can shed light on the structures of jets and their time evolutions. The variability mechanisms can be both intrinsic and extrinsic.



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

⁹ <https://www.cosmos.esa.int/web/euclid/euclid-survey>

¹⁰ In some objects, up to the X-ray band.

The intrinsic mechanisms are probably due to particle injection or acceleration in the jet, possibly produced by shock waves propagating downstream, or to magnetic reconnection (e.g., Böttcher & Baring 2019; Bodo et al. 2021). The extrinsic mechanisms have a geometrical nature, caused by changes in the orientation of the jet-emitting regions with respect to the line of sight, which produce variations in the corresponding Doppler beaming factor (e.g., Villata & Raiteri 1999; Raiteri et al. 2017). Both intrinsic and extrinsic mechanisms may be at work on different timescales.

If multiband optical light curves are available, as in Rubin-LSST, we can follow the spectral changes by means of color indices, provided these are obtained with data taken close in time. In the case of nearly contemporaneous data at several different frequencies in the optical band and beyond, we can build broadband spectral energy distributions (SEDs). The shape of the broadband SED can identify the blazar type—i.e. whether it is low-, intermediate-, or high-energy peaked, depending on the frequency ν of the synchrotron peak in the $\log(\nu F_\nu)$ versus $\log \nu$ representation, where F_ν is the flux density. Some sources have been seen to change type according to their brightness level. The spectral behavior over time depends on the variability mechanism: achromatism or quasi-achromatism is compatible with a geometrical interpretation, while strong chromatism implies energetic processes. In some cases, spectral changes are found to be much stronger on short than on long timescales. This has been interpreted as the result of rapid, intrinsic variability in jet regions that undergo orientation changes (Raiteri et al. 2021a, 2021b).

In this paper, we will investigate the impacts of the different choices for the Rubin-LSST observing strategy on blazar variability studies. We will make use of the well-sampled and carefully checked blazar light curves of the Whole Earth Blazar Telescope¹¹ (WEBT; e.g., Villata et al. 2002; Raiteri et al. 2017). Established in 1997, the WEBT is an international collaboration between many tens of astronomers, who monitor a number of sources—mainly in the optical band, but also at radio and, less frequently, at near-infrared wavelengths. We will adopt Roma-BZCAT5 (Massaro et al. 2009, 2015) as the reference blazar catalog, since it includes confirmed sources and estimates of optical brightness. However, we will also discuss the implications of other blazar catalogs containing tens of thousands candidates.

2. Emission Contributions in the Optical Band

The blazar AGN class includes flat-spectrum radio quasars (FSRQs) and BL Lac-type objects. They are classically distinguished on the basis of the strength of their emission lines: BL Lac-type sources show featureless spectra or lines with equivalent widths smaller than 5 Å in the rest frame (Stickel et al. 1991; Stocke et al. 1991). This definition is somewhat problematic, because it depends on the relative importance of the featureless jet emission contribution as opposed to that coming from the nuclear region, which includes the accretion disk and the broadline and narrowline regions. Even the prototype for BL Lac-type objects, BL Lacertae, does not behave as a BL Lac-type source when its synchrotron emission is weak (Vermeulen et al. 1995; Corbett et al. 2000). While in the radio and high-energy bands the beamed nonthermal radiation from the jet dominates, in the optical

the scenario is more complex, and we can generally encounter three different situations, which are described below.

2.1. Host-dominated BL Lac-type Objects

The optical emission of close BL Lac-type sources can be dominated by the light of the host galaxy. In BZCAT5, these objects represent $\sim 7.7\%$ of the whole blazar population, with the most distant one being 5BZGJ1552 + 3159 at redshift $z = 0.584$. These low-luminosity sources are recognized as blazars because of their multiwavelength behaviors and/or because of anomalous Ca break values (e.g., Capetti & Raiteri 2015). The study of the flux and spectral variability of the jet component requires the determination of the host galaxy's contribution to the total emission in the various bands, in order to subtract the host galaxy's light from the total flux.

2.2. Disk-dominated Flat-spectrum Radio Quasars

When the thermal radiation due to the accretion disk (the “big blue bump”) prevails over the nonthermal jet emission, these objects behave like normal quasars and their variability is well represented as a damped random walk (e.g., MacLeod et al. 2010; Butler & Bloom 2011). If a strong nonthermal flare occurs, it can be distinguished from the disk variability as a more rapid event that is likely to be seen outside the optical band, particularly at γ -rays.

The spectra of these objects show strong broad emission lines that can be detected as excess in some specific Rubin-LSST bands, which depend on the redshift, and can likely allow photometric reverberation mapping studies (e.g., Chelouche & Daniel 2012).

2.3. Synchrotron-dominated Sources

The most interesting sources from the point of view of blazar variability and jet understanding are BL Lac-type objects and FSRQs where nonthermal radiation plays the major role. The observed emission is subject to the rules of relativistic beaming, with Doppler-enhanced flux and variability amplitudes, and Doppler-shortened timescales. The most extreme objects are called optically violent variable (OVV) sources. Long-term and well-sampled optical light curves in different bands are needed to cover the variety of variability timescales. Spectral variability prevents the use of data in different filters to increase the sampling of a given light curve: if average color indices are adopted, the corresponding errors may reach several tenths of mag.

These sources can undergo major outbursts, with increases in brightness of up to several magnitudes (e.g., Raiteri et al. 2017), meaning that saturation in the Rubin-LSST data will affect some of the most critical and interesting events. In the following sections, we will address this problem.

3. Rubin-LSST Observing Strategies

The observing strategy of Rubin-LSST must meet the original goals (Ivezić et al. 2019) and, at the same time, optimize as many science cases as possible. Many cadence simulations (OpSim runs) have been performed to allow researchers in all fields of astronomy to test their science cases through the Metric Analysis Framework (MAF; Jones et al. 2014) software. These simulations cover many different possibilities, like changing the survey footprint or filter

¹¹ <http://www.oato.inaf.it/blazars/webt>




Figure 1. Map of the sky with the blazars in the BZCAT5, BROS, and CRATES catalogs plotted with black, gray, and green dots, respectively. BZCAT5 objects with R -band catalog magnitudes brighter than 15.5, which is the median maximum saturation magnitude in the LSST r band, are highlighted with blue circles. The regions enclosed by red lines represent the planned DDFs—from north to south, these are: COSMOS, XMM-LSS, ECDFS, ELAIS S1, and EDFs (the double field).

distribution, adding a third visit to the standard two visits per field in the same night, considering rolling cadences (with some regions of the sky being better sampled at certain periods of time), and adding short exposures a few times per year. The international astronomical community has been requested to provide metrics to test the performance of the various OpSim runs in relation to the scientific goals (Bianco et al. 2022). A detailed description and analysis of the cadence simulations can be found in Jones et al. (2020).¹²

In this paper, we will use metrics to explore how different OpSim runs impact the investigation of blazar variability. We will adopt the OpSim run `baseline_nexp2_v1.7_10yrs` as the reference cadence simulation. This includes two consecutive snapshots with exposure times of 15 s each for both of the two visits to a given field in the same night.

4. How Many Blazars?

Figure 1 shows the distribution in the sky of the 3561 blazars listed in the BZCAT5 catalog.¹³ The objects of BZCAT5 were collected from the literature and selected in different ways. The catalog gives multifrequency information, including optical

magnitudes in the R band down to 24.4 mag. In the figure, the locations of the selected DDFs are indicated, in accordance with Jones et al. (2020). In order of decreasing decl., we find: COSMOS, XMM-LSS, ECDFS, ELAIS S1, and EDFs (the double field matching the Euclid Deep Field South). Three of the DDFs do not contain BZCAT5 blazars. The other three DDFs only include five sources overall.

Were BZCAT5 a complete catalog, the DDFs would be of no interest for blazar studies. However, there are other catalogs of blazar candidates that present many more objects. The CRATES catalog (Healey et al. 2007) lists about 11,000 flat-spectrum radio sources outside the Galactic Plane. The largest catalog of blazar candidates is BROS (Itoh et al. 2020), which includes 88,211 sources with flat radio spectra at decl. $\delta > -40^\circ$ outside the Galactic Plane ($|b| > 10^\circ$). Optical magnitudes are available in the g , r , and i bands for nearly half of the sample. Of the 32,980 BROS sources with r -band magnitudes, only 12 sources exceed $r = 24.5$, the minimum single-visit depth of Rubin-LSST (Ivezić et al. 2019). The positions in the sky of the objects of the BROS and CRATES catalogs are shown in Figure 1. Table 1 displays the number of sources found by cross-matching the BZCAT5, BROS, and CRATES catalogs with the DDF positions. The total number of candidates in the final column only takes into account nonduplicated sources. The EDFs double DDF has been split into two separate components.

¹² Lists of the OpSim runs included in the v1.5, v1.6, and v1.7 releases are available at <http://astro-lsst-01.astro.washington.edu:8081>, where the results of the application of standard metrics are also given.

¹³ <https://www.asdc.asi.it/bzcat> (Massaro et al. 2009, 2015).

Table 1

Number of Blazar Candidates in Each DDF (in Decreasing Decl. Order),
According to the BZCAT5, BROS, and CRATES Catalogs

DDF (1)	BZCAT5 (2)	BROS (3)	CRATES (4)	N_{tot} (5)
COSMOS	1	25	1	26
XMM-LSS	0	46	2	46
ECDFS	0	29	4	31
ELAIS	0	0	2	2
EDFS	2	0	9	9
EDFS	2	0	3	3
Total				117

Note. The final column reports the total number of nonduplicated sources.

In total, there are 117 blazar candidates in the planned DDFs according to the above three catalogs. Most of them are faint objects, down to $r \sim 23.1$; only two sources are brighter than 15.5 mag, the median r -band saturation limit. If we consider that three DDFs are outside the sky area covered by BROS, and the depth that Rubin-LSST is able to reach, the overall number of expected blazar candidates in the DDFs further increases. The possibility of following these objects with more sampling—in terms of both light curves and colors—than in the WFD, and with limited saturation issues, makes them a small golden sample for investigation, with important scientific returns.

Taking into account the BROS sky coverage, we can estimate an average number density of about 3.30 sources per square degree, which would imply $\sim 136,000$ blazar “candidates” in the whole sky, with about 66,000 being in the $\sim 20,000$ square degrees covered by Rubin-LSST and ~ 190 in the six DDFs. We can then wonder whether this number makes sense: in other words, does the blazar population that will be seen by Rubin-LSST reasonably constitute several tens of thousands of objects? In order to make a rough estimate, we can consider that the quasar population to be explored by Rubin-LSST is expected to include about 10^7 objects up to $z \sim 7$, and that multiband light curves will be available to study quasar variability for a few million of these (Ivezić 2017; Ivezić et al. 2019). A fraction ζ of these quasars will be radio-loud. If one of their two jets were to be within $\sim 30^\circ$ of the line of sight, we would see a blazar. Therefore, we can obtain an estimate of the number of blazars that Rubin-LSST is likely to observe by scaling the number of radio-loud quasars by the ratio between the solid angle corresponding to a 60° aperture and that of the half sphere (which is 2π):

$$N \approx 10^7 \times \zeta \times 2\pi(1 - \cos(\pi/6))/(2\pi). \quad (1)$$

The radio-loud fraction ζ is commonly assumed to be around 10% (e.g., Ivezić et al. 2002), but it has actually been found to decrease from about 24% to roughly 4% with increasing redshift and decreasing luminosity (Jiang et al. 2007; Kratzer & Richards 2015). By adopting the latter value, we would obtain a lower limit of about 53,600 objects, which is of the order of the BROS extrapolation. By considering a Rubin-LSST footprint of $\sim 20,000$ square degrees, we would obtain ~ 2.68 sources per square degree, and about 150 sources in the six planned DDFs, most of which will presumably be detected with a single-visit exposure.

5. Saturation

The Rubin-LSST photometry of the brightest blazars, or blazars during outburst states, can be affected by saturation. In photometric conditions with $0.7''$ seeing, the saturation levels for a 15 s exposure are expected to be 14.7, 15.7, 15.8, 15.8, 15.3, and 13.9 mag for the u , g , r , i , z , and y bands, respectively.¹⁴ These values increase with the exposure time, but since the saturation levels change with the observing conditions, they can vary significantly over time and with the position in the sky. Figure 2 shows the median (over the sky) saturation magnitudes obtained with both the OpSim run `baseline_nexp2_v1.7_10yrs`—which includes two snapshots with 15 s exposures for all visits—and with the `baseline_nexpl_v1.7_10yrs` cadence, where visits consist of single exposures of 30 s. For each band, the maximum, median, and minimum saturation magnitudes obtained from the simulated 10 yr period of the survey (which includes variable observing conditions) are given. In the 30 s single exposure case, the level of saturation increases by about 0.7–0.8 mag with respect to the 15 s double exposures. For a better comprehension of these values, Figure 2 also includes a sky map showing the faintest objects that could saturate in the r band, and the corresponding histogram, for the `nexp2` cadence simulation.

To investigate the problem of saturation, we can limit ourselves to the BZCAT5 catalog, in which the brightest and more active sources are probably all included, and which offers the advantage of providing values for optical brightness. Although the catalog magnitude for a variable object represents just one of its possible brightness states, we will use these magnitudes as reference values, with the idea that they cover the whole range of brightness states of the known blazar population. The 191 sources with R -band catalog magnitudes brighter than 15.5, which represents the median maximum saturation level in the Rubin-LSST r band, are highlighted in Figure 1. They would saturate in 15 s exposures, when observed in the best seeing conditions.

We performed a simulation (`BlazarSaturation-Metric`) where all of the blazars from BZCAT5 are flaring according to given prescriptions. They go “on-flare,” starting from the catalog magnitudes, which are considered to be the “off-flare” states. The flare amplitudes, temporal distances between the flares, and flare durations are randomly chosen from a normal distribution of a range of reasonable values. Figure 3 shows the distribution of: (i) the blazar catalog magnitudes used as the “off-flare” brightness levels; (ii) the flare amplitudes, corresponding to the (maximum–minimum) magnitudes; (iii) the time intervals (in days) for which the sources remain “on-flare” (i.e. the plateau time); (iv) and the time intervals (in days) between consecutive flares (period). The same figure also shows the percentage of sources detected—100% would be reached if Rubin-LSST could see the whole sky—and the percentage of sources detected in outburst (i.e., that are observed to undergo $\Delta\text{mag} > 0.5$ with good photometry). Finally, the figure also displays the percentage of saturated observations for the various OpSim runs (see Section 3).

The largest numbers of detected sources and sources detected in outburst are obtained in the runs identified as `alt_dust`,

¹⁴ From the Rubin-LSST Science Book: <https://www.lsst.org/scientists/scibook> (LSST Science Collaboration 2009).




Figure 2. (a) Median values over the sky of the maximum (blue squares), median (black dots), and minimum (red diamonds) saturation magnitudes over the 10 yr period of the survey for the different Rubin-LSST filters. The filled symbols refer to the OpSim run `baseline_nexp2_v1.7_10yrs`, which include two exposures of 15 s each per visit; the unfilled symbols refer to `baseline_nexp1_v1.7_10yrs`, with a single exposure of 30 s per visit. Histogram (b) and sky map (c) of the maximum saturation magnitude in the *r* band for the `nexp2` cadence simulation.

`alt_roll_mod2_dust`, `footprint_big_sky_dust`, and `roll_mod2_dust`. All of these simulations imply an extension of the survey footprint.

The number of saturated observations goes from $\sim 3\%$ to $\sim 6\%$, and the OpSim runs with single exposures of 30 s perform worse than the runs with two exposures of 15 s, as expected. Intermediate saturation levels are obtained with the `short_exp` cadences, which include a minisurvey with either 1 or 5 s exposures, either twice or five times per year, and also by the `filterdist` cadences, where the filter distribution is changed with respect to the `baseline` simulation.

In Figure 4, we plot sky maps highlighting whether a source has been observed to flare (left) and the fraction of saturated observations (right), as obtained with the `baseline_nexp2_v1.7_10yrs` OpSim run.

We stress that although such saturation seems to affect only a small fraction of the outbursts, these bright events are the most important ones for understanding jet physics, especially in a multiwavelength context. Saturation would hamper the possibility of triggering follow-up observations, and thus exploiting

the synergy between Rubin-LSST and high-energy facilities in order to understand the role played by leptonic versus hadronic processes in producing the X-ray and γ -ray radiation in blazars, and also, in turn, the emission of high-energy neutrinos.

6. Flux and Color Sampling

6.1. Baseline Cadence

The study of blazar variability requires well-sampled light curves and color indices. We first explore the number of data points that can be obtained by both the WFD and the DDFs with the `baseline_nexp2_v1.7_10yrs` cadence simulation over 10 yr. We also investigate the number of color indices that can be built from data acquired the same night in order to mitigate any possible bias introduced by variability. The results are shown in Figure 5 for both the DDFs and the WFD. In the latter case, we tested three different sky positions. The numbers are brightness-agnostic; they do not depend on the light curve features (provided that all observations lead to a detection) and they do not take saturation into account. The light curve




Figure 3. Left: parameter distributions for the flaring sources used in the saturation metric. From top left, clockwise: the *R*-band catalog magnitudes adopted for “off-flare” brightness levels; the outburst amplitudes (mag); the plateau times (days); and the periods (days). Right: the top panel shows the percentage of detected sources (red) and the percentage of sources detected in outburst (blue), while the bottom panel shows the percentage of saturated observations for the various OpSim runs listed on the right.




Figure 4. Left: sky map of 2029 simulated blazars. Values of 1 denote they would be detected brightening by 0.5 mag in the baseline strategy, values of zero mean there would be no detection. Right: the fraction of observations per source that would be saturated.




Figure 5. The number of observations (left) and the number of color indices obtained by coupling the data from two bands taken in the same night (right) at the end of the 10 yr survey. Values have been obtained using the baseline_nexp2_v1.7_10yrs OpSim run. The black triangles refer to the DDFs, while the other symbols and colors refer to the WFD, in three different positions in the sky.




Figure 6. The minimum time separations (in days) between the observations in two different filters for the WFD (left) and the DDFs (right). Values have been obtained using the `baseline_nexp2_v1.7_10yrs` OpSim run. Different colors are used for different filter pairs, and the horizontal width of each pair is proportional to the corresponding number of points, in order to better distinguish the number density. The black crosses indicate mean values.

sampling is best in the r , i , z , and y filters in the WFD, and in the r , i , and y bands in the DDFs. The number of data points in the DDFs can be more than 20 times greater than those in the WFD. The best-sampled colors are: $g-r$, $r-i$, $i-z$, and $z-y$ for the WFD, and $r-i$, $r-y$, and $i-y$ for the DDFs. The maximum number of color index measurements with data for the same night in the WFD is less than 100 over 10 yr, while in the DDFs we have more than 3500 indices in the best cases.

We can wonder by how much we should increase the time separations between the observations in two different filters either to significantly improve the number of corresponding color index measurements—and thus trace the spectral changes more efficiently—or even to obtain colors that cannot be obtained with the same-day constraint. Figure 6 displays the distribution of the minimum time separations between acquisitions of data in the two different bands. Results for both the WFD and the DDFs are shown. The choice of the time separation limit should depend on how fast the variability of the source is; for OVV sources, the same-day constraint seems a necessary requirement, but for the other (slowly variable) objects, a time separation of a few days may be acceptable.

For the WFD, the `baseline` cadence implies a median value (over the sky) of 3.83 days for the median (over 10 yr) internight time gaps between consecutive observations of the same field, when all bands are considered. When single bands are distinguished, the median time gaps are 21.94, 18.46, 7.87, 9.02, 11.89, and 20.85 days in the u , g , r , i , z , and y bands, respectively.

To see the effects of Rubin-LSST sampling in a real case, we used the processed $BVRI$ light curves of OVV 3C 454.3, which were obtained by the monitoring efforts of the WEBT and published in Villata et al. (2006, 2007, 2009) and Raiteri et al. (2007, 2008a, 2008b, 2011). These are shown in Figure 7, and cover a time interval of 1980 days, i.e., ~ 5.425 yr. We have adopted the `baseline_nexp2_v1.7_10yrs` cadence and the `TransitAsciiMetric` available in the MAF. Since the code makes interpolations between data points, and the sampling of the WEBT light curves is not the same for all bands, we have filled the gaps in the light curves in a controlled way in order to avoid different interpolated trends in different bands. To do this, we first ran a cubic spline interpolation on the 4 day binned R -band light curve, which is the best-sampled one. We then calculated the average color indices by coupling the data within 30 minutes. We got: $\langle B-R \rangle = 1.13$, $\langle V-R \rangle = 0.47$, and $\langle R-I \rangle = 0.62$. These

values were used to shift the spline to match the light curves in the other bands. As input to the `TransitAsciiMetric`, we gave the $BVRI$ light curves containing WEBT data with errors less than 0.1 mag, complemented by spline data (shifted according to the corresponding mean color index) for the days when real data were not available (see Figure 7). Since the most-sampled Rubin-LSST light curves are in the r , i , z , and y bands, we established the correspondence $B=r$, $V=i$, $R=z$, and $I=y$. Because the cadence simulation is run for the whole length of the survey, at the end of the period covered by the WEBT observations, the code repeats the light curves from the beginning in order to cover the entire 10 yr period.

The simulated light curves obtained by running the `baseline` OpSim on the 3C 454.3 interpolated WEBT light curves are plotted in Figure 7, where the sampling of both the WFD and the DDFs are also considered. We note that the counterpart to the peak of the outstanding outburst from 2005 in the 10 yr simulation would go undetected by Rubin-LSST in 2023, but that it would be fairly sampled in 2028, even by the WFD. Table 2 reports the number of data points in the simulated $rizy$ light curves of Figure 7, as well as the corresponding number of observing nights, to highlight the intranight sampling. In the WFD, the average number of points per night is 1.3–1.4 in the r , i , and z bands, and it is 2.1 in the y band, implying more intranight observations in this filter. These values rise to 10–14 data points per night in the DDFs. Therefore, if the number of data points in the DDFs is 11–20 times greater than those in the WFD, the observing nights are only 1.4–3 times more numerous. This is why we cannot see much better sampling in the DDFs with respect to the sampling in the WFD in Figure 7. Moreover, the total duration of the intranight observing sequence in the same filter is about 11 minutes, which is generally too short to follow even the fastest flux changes of OVV blazars. The distribution of the time separations between subsequent DDF observations is shown in Figure 8: about 90% of the cases lie in the first bin, indicating a gap of less than 1 minute. Table 2 and Figure 8 also show a different choice for the DDF cadence, which will be discussed in Section 6.2.

As shown in Figure 5, the color indices with better sampling in the WFD are $r-i$, $i-z$, and $z-y$. This means $B-V$, $V-R$, and $R-I$ for the WEBT data, according to the correspondence we set. Figure 9 shows the color indices obtained from the WEBT data with a time difference between the two filters of 1 hr. As mentioned before, a small time interval is needed because of




Figure 7. Left: the $BVRi$ light curves of the blazar 3C 454.3 processed by the WEBT. The red line in the R -band panel represents a cubic spline interpolation through the 4 day binned light curve. The red lines in the other panels represent the same cubic spline interpolation, shifted according to average color indices. Right: the simulated Rubin-LSST sampling, according to `TransitAsciiMetric`, with a 10 yr baseline cadence. The blue diamonds refer to the WFD, while the green plus signs refer to the DDFs. The gray symbols indicate the WEBT data shown in the left panel.

Table 2

Light-curve Sampling in Terms of Number of Data Points and Observing Nights in the Case of the `baseline_nexp2_v1.7_10yrs` Cadence Simulation for Both the WFD and the DDFs

Band (1)	WFD (baseline)		DDF (baseline)		DDF (daily)	
	N_{points} (2)	N_{nights} (3)	N_{points} (4)	N_{nights} (5)	N_{points} (6)	N_{nights} (7)
r	195	147	3740	298	3368	1058
i	191	149	3765	305	4323	1065
z	171	122	1831	176	2173	656
y	184	88	3699	270	4274	1048

Note. For the DDFs, we also show the results of the `daily_ddf_v1.5_10yrs` cadence.

the strong (and chromatic) intranight variability of OVV sources like 3C 454.3. In the considered period of 5.4 yr, utilizing the WEBT observations, we calculated 1147 $B-V$, 1672 $V-R$, and 1699 $R-I$ indices, with mean time separations between the observations in the two bands of about 7, 10, and 9 minutes, respectively. These are compared with the color indices obtained using the `baseline` cadence simulation on the interpolated WEBT light curves and the same time separation of 1 hr after 10 yr. For the WFD, we got 72 $r-i$,

99 $i-z$, and 47 $z-y$ indices, with mean time separations between the data in the two bands of 27, 24, and 26 minutes, respectively. If we relaxed the time separation to 1 day, the number of indices would become 82, 112, and 73, respectively. The WEBT and simulated WFD color indices are shown in Figure 9 as a function of both brightness and time. The same $\Delta \text{index} = 1.2 \text{ mag}$ in all panels makes a comparison between them easier.

When the two data points yielding the simulated color index are both taken from the spline interpolation, the resulting color index is close to the average value used to shift the cubic spline interpolation through the R -band data points to match the data in the other bands. This explains the accumulation of data points around the average color index. In any case, it seems that it is hard to sample the spectral changes in the brightest phases, and several fast color variations are lost. After a few years, however, the WFD would be able to assess the main features of blazar spectral variability: the fast variations appear to be more chromatic than the long-term ones, and there is a trend with brightness. In the case of 3C 454.3, the trend is “redder-when-brighter” above a certain magnitude ($R \gtrsim 14$), because of the presence of nuclear emission, which is steadier and bluer than the more variable synchrotron radiation. For increasing brightness, a kind of “saturation” effect appears, as discussed in Villata et al. (2006). In the figure, linear and parabolic fits have




Figure 8. Time separations (minutes) between the subsequent observations of the DDFs in the same filter. The different colors distinguish different bands, as indicated in the legend. The left panel refers to the `baseline_nexp2_v1.7_10yrs` cadence simulation, and the right panel refers to the `daily_ddf_v1.5_10yrs` one.




Figure 9. Color indices as a function of brightness (left) and time (right). The gray circles and blue diamonds represent real WEBT and simulated WFD colors, respectively. All of them have been obtained by coupling data with time separations of less than 1 hr. The red dashed line indicates the mean color index derived from the WEBT data. The accumulation of blue points close to the average level is explained in the text. The black lines are linear and parabolic fits to the WEBT data, highlighting the redder-when-brighter trend, with the “saturation” effect at the bright end. All plots have the same amplitude $\Delta \text{index} = 1.2 \text{ mag}$ to make the comparison between the colors easier. A few outliers, deviating more than 3σ from the mean, have been discarded.

been drawn in order to highlight this behavior, which can be better appreciated when the color involves redder filters, including more variable synchrotron contribution.

In the DDF `baseline` cadence simulation, the best color sampling is reached for $r-i$, $r-y$, and $i-y$, which correspond to

$B-V$, $B-I$, and $V-I$. As can be seen from Table 3, the numbers of $r-i$, $r-y$, and $i-y$ color indices that can be obtained in the DDFs are 50, 36, and 75 times greater than those in the WFD, respectively. However, the numbers of nights on which they are obtained are only 4.4, 2.3, and 4.5 times greater. As in the




Figure 10. Color indices as a function of brightness (left) and time (right). The gray circles and green plus signs represent the real WEBT and simulated DDF colors, respectively. All of them have been obtained by coupling data with time separations below 1 hr. The red dashed line indicates the mean color index derived from the WEBT data. The accumulation of green points close to the average level is explained in the text. The black lines are linear and parabolic fits to the WEBT data, highlighting the redder-when-brighter trend, with the “saturation” effect at the bright end. All plots have the same amplitude Δ index = 1.2 mag to make the comparison between the colors easier. A few outliers, deviating more than 3σ from the mean, have been discarded.

Table 3

Color Indices Sampling in Terms of Number of Data Points and Observing Nights in the Case of the `baseline_nexp2_v1.7_10yrs` Cadence Simulation for Both the WFD and the DDFs

WFD			DDF		
Index (1)	N_{points} (2)	N_{nights} (3)	Index (4)	N_{points} (5)	N_{nights} (6)
$r-i$	72	51	$r-i$	3634	225
$i-z$	99	76	$r-y$	3531	177
$z-y$	47	40	$i-y$	3534	180

case of flux, the duration of the intranight sampling for colors is also too short to follow the spectral variability even of OVV sources. The $r-i$, $r-y$, and $i-y$ color indices obtained with the WEBT data by coupling observations within 1 hr are displayed in Figure 10, where they are compared to those obtained with the baseline DDF cadence simulation. The same Δ index = 1.2 mag of Figure 9 has also been adopted here, for an easier comparison between all of the color behaviors. We note that both the amplitudes of the color variability and the color trends with brightness are more noticeable when the two bands are further apart. Indeed, while the $r-i$ color indices only cover a small range of values, and the trend with brightness is weak, the $i-y$ colors show a larger range and a more defined

trend, and the $r-y$ colors span almost the whole Δ index range, with linear and parabolic fits of the color-versus-brightness plot indicating a larger slope and a more pronounced curvature than in the other cases.

As mentioned in the introduction, the color variability shown in Figures 9 and 10 is a property of the source, and its interpretation can shed light on the contributions to the source emission and on the physical conditions in the jet. Therefore, it would also be present in the case of simultaneous photometry in two bands, but may be altered by flux variability, if the observations in the two bands are too distant in terms of time.

6.2. Other Cadences

Until now, we have explored light curve and color sampling according to the `baseline_nexp2_v1.7_10yrs` cadence. In this section, we investigate whether other OpSim runs can lead to better results. Figure 11 reports the difference between the number of observations obtained over 10 yr using various WFD cadence simulations, with respect to the `baseline_nexp2_v1.7_10yrs` OpSim run. We show cadences including two 15 s exposures per visit and cadences with single 30 s exposures. Of the former, none seem more favorable than the baseline simulation, overall. The single 30 s exposure cadences lead to some time saving and so to more observations. Apart from the




Figure 11. Difference between the numbers of WFD visits in various cadence simulations and that of the `baseline_nexp2_v1.7_10yrs` cadence. OpSim runs with either double 15 s (left) or single 30 s (right) exposure times per visit are shown.




Figure 12. Difference between the numbers of DDF visits in various cadence simulations and that of the `baseline_nexp2_v1.7_10yrs` cadence. OpSim runs with either double 15 s (left) or single 30 s (right) exposure times are shown.

two extreme `filterdist` simulations, which strongly favor either blue or red filters, with supersampling in these filters and undersampling in the other bands, the other OpSim runs can provide more visits than the baseline in some filters, but in general the difference is much less than 50 visits over 10 yr. This small advantage would occur at the expense of doubling the number of saturated observations. In particular, the `baseline_nexp1_v1.7_10yrs` yields -2 , -3 , 10 , 20 , 18 , and -11 data points in the u , g , r , i , z , and y bands over 10 yr with respect to the `baseline_nexp2_v1.7_10yrs`, which is a negligible difference.

The situation changes dramatically in the case of the DDFs, as shown in Figure 12. Some of the rolling cadence simulations with two 15 s exposures imply many more observations than the corresponding baseline cadence. The number further increases with single 30 s exposures. The OpSim run `baseline_nexp1_v1.7_10yrs` provides 105, 187, 365, 340, 159, and 387 more data points in the u , g , r , i , z , and y filters, respectively. The `rolling_extragalactic` cadence appears as an extreme case, yielding 86, 292, 603, 582, 483, and 587 more visits. As detailed by Jones et al. (2020), this cadence concentrates on low-extinction regions and divides the sky into quarters, so there is always a region of observing emphasis that is reachable by northern telescopes. However, the increase of 603 data points in the r band corresponds to an increase of only 22 nights over 10 yr,

which is not a great gain, especially when considering that it would worsen the saturation problem.

There is actually a series of OpSim runs that specifically test different DDF cadences (Jones et al. 2020). In particular, the `daily_ddf_v1.5_10yrs` observing strategy leads to a much larger number of nights, as can be seen in Table 2. However, even in this case, the interday and intraday samplings are not optimal. This is shown in Figure 8, where the time difference between subsequent observations in the daily cadence can be compared to that in the baseline cadence. Although the daily cadence leads to a better sampling of the interday timescales, and to a slightly better sampling of the intraday timescales, a large percentage of all of the observation pairs (72% for the i and y bands, and 64% for the r and z bands) still have time separations within 1 minute.

7. Discussion and Conclusions

In this paper, we have analyzed the impacts of different choices for the Rubin-LSST observing strategy on the study of both brightness and spectral blazar variability. We first discussed the number of blazar candidates that Rubin-LSST is likely to see, and found that it is of the order of several tens of thousands. Most of them should be detectable with a single-visit exposure, making a multiband variability study possible. This estimate was obtained both from general considerations and by taking the largest catalog

of blazar candidates into account. If it is correct, besides the WFD, the DDFs will be of interest for blazar variability studies, as they will contain more than a hundred objects for which more detailed investigations of flux and spectral variability will be possible.

Saturation is a serious problem, since it could affect the most interesting events (outbursts) that are eligible for multiwavelength follow-up observations. Saturation will also affect the brightest blazars, many of which are alleged to be sources of high-energy neutrinos. In the white paper by Raiteri et al. (2018), we suggested a minisurvey in star trail mode, which could solve the problem of blazar saturation, and at the same time offer the opportunity to search for very rapid transients, such as possible optical counterparts of fast radio bursts. We are aware that this idea may present technical problems and may be expensive in terms of time. Another possibility for overcoming the saturation problem, which was proposed in the same paper, would be to diversify the two exposure times of a single visit, using, e.g., $2 + 28$ s or $5 + 25$ s snapshots. In the available OpSim runs, the number of saturated observations in a simulation with reasonable flaring conditions goes from about 3% to about 6% when we compare cadences with two exposures of 15 s per visit to those with single 30 s exposures. Some cadences with single 30 s exposures provide many more observations, especially in the DDFs, but only a few more observing nights. Indeed, DDFs are characterized by sequences of intranight observations in the same filter, but these sequences are so short—about 11 minutes—that they cannot generally trace the fastest flux and spectral changes, even in OVV sources.

As mentioned, the understanding of blazar variability requires intensive multiwavelength monitoring. In the case of outstanding flaring events, follow-up observations should be triggered, especially at X-ray and γ -ray energies. Indeed, the high-energy counterparts of optical flares usually occur almost contemporaneously, meaning that prompt alerts are required. This can be accomplished through the Rubin-LSST brokers (Bellm et al. 2020), which are expected to release millions of alerts per night for objects that can be detected with at least 5σ on the difference images. Many of these alerts will concern blazars. Beneficial follow-up observations should also be carried out in polarimetric mode. Indeed, blazar synchrotron emission is polarized, and both the polarization degree and angle can be quite variable. The polarimetric behavior not only gives information about the magnetic field in the jet, but also about the jet's geometrical structure (e.g., Larionov et al. 2013; Raiteri & Villata 2021).

Our recommendations for the choice of the Rubin-LSST cadence from the point of view of blazar variability are:

1. In the WFD, double 15 s exposures are favored with respect to single 30 s ones, in order to mitigate the saturation of flaring/bright blazars. No cadence was found to perform much better than the OpSim run `baseline_nexp2_v1.7_10yrs`, whose major drawback is the coupling of adjacent filters in the same night (see below).
2. In the DDFs, the problem of saturation should have a limited impact, and cadences with single 30 s exposures that provide more sampling can be considered. However, a more beneficial cadence should include shorter DDF exposure sequences more often; and, in particular, intranight observations of the same DDF with more time spacing between visits in the same filter. In this sense, a DDF observing strategy like that included in the `daily_ddf_v1.5_10yrs` OpSim run represents a better choice than the `baseline` implementation, but this can be further improved by introducing a more homogeneous sampling of

the different variability timescales (see also Bellm et al. 2021).

3. In both the WFD and the DDFs, visits with different filters in the same night are required in order to obtain color indices with data taken close in time, which will be necessary to avoid bias in colors due to variability; this means that the `samefilt` OpSim runs are detrimental.
4. Color variability and trends with brightness are clearer when colors are obtained with data in filters that are further apart by wavelength; therefore, in both the WFD and the DDFs, the choice of filters for coupling in the same night should prefer bands that are not contiguous.
5. We favor an extension of the WFD footprint to the north, with an enlargement of the low-extinction extragalactic sky, which would also be reachable by many more observing facilities that could complement the Rubin-LSST monitoring. This should be accomplished without decreasing the sampling.
6. We cannot support a distribution of filters more skewed toward blue/red filters than in the `baseline` cadence, because different types of blazars would require different choices of filters.
7. Observations at high air masses may be beneficial, as they can prolong the observing season and lead to smaller gaps in the light curves. Moreover, high air mass can also avoid saturation in some cases.
8. Target of opportunity (ToO) observations should be considered when high-energy neutrinos of astrophysical origin are detected by neutrino facilities. The uncertainty regarding the arrival direction can be up to a number of degrees. Since the most promising candidates for neutrino sources are bright blazars, the ToO observations should include short exposures to avoid possible saturation.

We finally mention that, besides variability, Rubin-LSST will improve other aspects of our knowledge of blazars, like census and environment. Rubin-LSST will potentially identify thousands of new objects. These will likely be FSRQs at high redshift, with important cosmological implications. As already mentioned, blazars belong to the radio-loud AGN population, which seems to require SMBH with masses greater than $10^8 M_\odot$. Blazars at high redshift thus imply the presence of SMBH in the early universe. Until now, the blazar with the highest redshift ($z > 6$) was discovered by Belladitta et al. (2020). The identification of new blazar candidates with Rubin-LSST must rely on a combination of colors and variability features, and possibly on multiwavelength information, also thanks to new observing facilities with enhanced capabilities. In the X-ray band, the extended ROentgen Survey with an Imaging Telescope Array (eROSITA; Predehl et al. 2021) has been performing an all-sky survey since 2019, with a sensitivity of $10^{-14} \text{ erg cm}^{-2} \text{ s}^{-1}$. The survey will include eight scans of the whole sky over 4 yr. The Advanced Telescope for High Energy Astrophysics (Athena; Barcons et al. 2017), to be launched in 2030, will perform various surveys, further improving depth. In particular, the “wide” survey will cover 48 deg^2 and include the Rubin-LSST DDFs, with a sensitivity of $\sim 10^{-16} \text{ erg cm}^{-2} \text{ s}^{-1}$ in the 0.5–2.0 keV energy range. For comparison, the 0.1–2.4 keV flux of the 2249 blazars in BZCAT5 for which X-ray data are available ranges between 2×10^{-14} and $3.2 \times 10^{-10} \text{ erg cm}^{-2} \text{ s}^{-1}$. An ultrawide ($\sim 800 \text{ deg}^2$) and shallow ($\sim 5 \times 10^{-16} \text{ erg cm}^{-2} \text{ s}^{-1}$) Athena survey, which would complement Rubin-LSST nicely, although still over a smaller area, is under consideration. At γ -rays, the Fermi satellite (Atwood et al. 2009) has been scanning the entire sky each

day since 2008. The Fermi Large Area Telescope (LAT) fourth source catalog (4FGL; Abdollahi et al. 2020) includes 5064 sources, $\sim 62\%$ of which are blazars, but there are still 1336 unassociated sources, most of which are expected to be blazars, for which Rubin-LSST may find the optical counterparts. Finally, 80 blazars are detected at very high energies ($E > 100$ GeV) by current Cherenkov telescopes.¹⁵ 50 of them lie at $-90^\circ < \delta < +30^\circ$ and are likely to be observable by Rubin-LSST. According to BZCAT5, their magnitude is $11 < R < 19.4$, with 25 sources brighter than 15.5 mag, and thus possibly affected by saturation. Moreover, the second catalog of hard Fermi-LAT sources (2FHL; Ackermann et al. 2016) reports the detection of 265 blazars in the 50 GeV–2 TeV energy range, most of them potentially detectable by the next-generation Cherenkov Telescope Array (Acharya et al. 2019), and whose optical counterparts are expected to be identifiable by Rubin-LSST. The blazar candidates will then need validation through follow-up spectroscopic observations.

Finally, the depth of the Rubin-LSST observations will allow us to study the blazar environment, which in turn is a fundamental ingredient for understanding what the parent population of blazars is among unbeamed AGN sources. Muriel (2016) found that a high percentage of BL Lac-type objects lie in groups of galaxies. However, Sandrinelli et al. (2019) claimed that the environment of BL Lac-type sources is less rich by a factor of 2 than that of Fanaroff-Riley type I (FR I) radio galaxies, which represent their parent population according to the AGN unification scheme. This result, which questions the identification of the parent population of the BL Lac-type objects, needs to be verified with larger and more uniform samples of FR I and BL Lac-type objects.

We are deeply grateful to Rachel Street, Federica Bianco, and Niel Brandt for their essential coordination work inside the Rubin-LSST Science Collaborations. This paper was created in the nursery of the Rubin-LSST “Transients and Variable Stars” Science Collaboration (<https://lsst-tvssc.github.io/>). The authors acknowledge the support of the Vera C. Rubin Legacy Survey of Space and Time “Transients and Variable Stars” and “AGN” Science Collaborations that provided opportunities for collaboration and the exchange of ideas and knowledge. The authors are thankful for the support provided by the Vera C. Rubin Observatory MAF team in the creation and implementation of MAFs. The authors acknowledge the support of the LSST Corporation, which enabled the organization of many workshops and hackathons throughout the cadence optimization process. This work was supported by the Preparing for Astrophysics with LSST Program, funded by the Heising Simons Foundation through grant 2021-2975, and administered by Las Cumbres Observatory.

The authors acknowledge the use of data taken and assembled by the WEBT collaboration and stored in the WEBT archive at the Osservatorio Astrofisico di Torino—INAF (<http://www.oato.inaf.it/blazars/webt/>).

Facility: Rubin.

Software: LSST metrics analysis framework (MAF; Jones et al. 2014).

ORCID iDs

Claudia M. Raiteri  <https://orcid.org/0000-0003-1784-2784>

Maria I. Carnerero  <https://orcid.org/0000-0001-5843-5515>

Barbara Balmaverde  <https://orcid.org/0000-0002-0690-0638>

Eric C. Bellm  <https://orcid.org/0000-0001-8018-5348>
 William Clarkson  <https://orcid.org/0000-0002-2577-8885>
 Filippo D’Ammando  <https://orcid.org/0000-0001-7618-7527>
 Maurizio Paolillo  <https://orcid.org/0000-0003-4210-7693>
 Gordon T. Richards  <https://orcid.org/0000-0002-1061-1804>
 Massimo Villata  <https://orcid.org/0000-0003-1743-6946>
 Peter Yoachim  <https://orcid.org/0000-0003-2874-6464>
 Ilsang Yoon  <https://orcid.org/0000-0001-9163-0064>

References

- Aartsen, M., Ackermann, M., Adams, J., et al. 2018a, *Sci*, **361**, eaat1378
 Aartsen, M., Ackermann, M., Adams, J., et al. 2018b, *Sci*, **361**, 147
 Abdollahi, S., Acero, F., Ackermann, M., et al. 2020, *ApJS*, **247**, 33
 Acharya, B. S., Agudo, I., Al Samarai, I., et al. 2019, *Science with the Cherenkov Telescope Array* (Singapore: World Scientific)
 Ackermann, M., Ajello, M., Atwood, W. B., et al. 2016, *ApJS*, **222**, 5
 Atwood, W. B., Abdo, A. A., Ackermann, M., et al. 2009, *ApJ*, **697**, 1071
 Barcons, X., Barret, D., Decourchelle, A., et al. 2017, *AN*, **338**, 153
 Belladitta, S., Moretti, A., Caccianiga, A., et al. 2020, *A&A*, **635**, L7
 Bellm, E., Blum, R., Graham, M., et al. 2020, LSST Document LDM-612, <https://ldm-612.lsst.io/>
 Bellm, E. C., Burke, C. J., Coughlin, M. W., et al. 2021, arXiv:2110.02314
 Bianco, F. B., Ivezić, Ž., Jones, R. L., et al. 2022, *ApJS*, **258**, 1
 Bodo, G., Tavecchio, F., & Sironi, L. 2021, *MNRAS*, **501**, 2836
 Böttcher, M., & Baring, M. G. 2019, *ApJ*, **887**, 133
 Böttcher, M., Reimer, A., Sweeney, K., & Prakash, A. 2013, *ApJ*, **768**, 54
 Butler, N. R., & Bloom, J. S. 2011, *AJ*, **141**, 93
 Capetti, A., & Raiteri, C. M. 2015, *A&A*, **580**, A73
 Chelouche, D., & Daniel, E. 2012, *ApJ*, **747**, 62
 Corbett, E. A., Robinson, A., Axon, D. J., & Hough, J. H. 2000, *MNRAS*, **311**, 485
 Giommi, P., Glauch, T., Padovani, P., et al. 2020, *MNRAS*, **497**, 865
 Healey, S. E., Romani, R. W., Taylor, G. B., et al. 2007, *ApJS*, **171**, 61
 Itoh, R., Utsumi, Y., Inoue, Y., et al. 2020, *ApJ*, **901**, 3
 Ivezić, Ž. 2017, in IAU Symp. 324, *New Frontiers in Black Hole Astrophysics* (Cambridge: Cambridge Univ. Press), 330
 Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111
 Ivezić, Ž., Menou, K., Knapp, G. R., et al. 2002, *AJ*, **124**, 2364
 Jiang, L., Fan, X., Ivezić, Ž., et al. 2007, *ApJ*, **656**, 680
 Jones, R. L., Yoachim, P., Chandrasekharan, S., et al. 2014, *Proc. SPIE*, **9149**, 91490B
 Jones, R. L., Yoachim, P., Ivezić, V., Neilsen, E. H., & Ribeiro, T. 2020, LSST Document PSTN-051 Zenodo, doi:10.5281/zenodo.4048838
 Kratzer, R. M., & Richards, G. T. 2015, *AJ*, **149**, 61
 Larionov, V. M., Jorstad, S. G., Marscher, A. P., et al. 2013, *ApJ*, **768**, 40
 LSST Science Collaboration 2009, *LSST Science Book*, v2.0., arXiv:0912.0201
 MacLeod, C. L., Ivezić, Ž., Kochanek, C. S., et al. 2010, *ApJ*, **721**, 1014
 Massaro, E., Giommi, P., Leto, C., et al. 2009, *A&A*, **495**, 691
 Massaro, E., Maselli, A., Leto, C., et al. 2015, *ApSS*, **357**, 75
 Murase, K., Dermer, C. D., Takami, H., & Migliori, G. 2012, *ApJ*, **749**, 63
 Muriel, H. 2016, *A&A*, **591**, L4
 Predehl, P., Andritschke, R., Arefiev, V., et al. 2021, *A&A*, **647**, A1
 Raiteri, C. M., Carnerero, M. I., Balmaverde, B., et al. 2018, arXiv:1812.03151
 Raiteri, C. M., & Villata, M. 2021, *Galax*, **9**, 42
 Raiteri, C. M., Villata, M., Acosta-Pulido, J. A., et al. 2017, *Natur*, **552**, 374
 Raiteri, C. M., Villata, M., Aller, M. F., et al. 2011, *A&A*, **534**, A87
 Raiteri, C. M., Villata, M., Carosati, D., et al. 2021a, *MNRAS*, **501**, 1100
 Raiteri, C. M., Villata, M., Chen, W. P., et al. 2008b, *A&A*, **485**, L17
 Raiteri, C. M., Villata, M., Larionov, V. M., et al. 2007, *A&A*, **473**, 819
 Raiteri, C. M., Villata, M., Larionov, V. M., et al. 2008a, *A&A*, **491**, 755
 Raiteri, C. M., Villata, M., Larionov, V. M., et al. 2021b, *MNRAS*, **504**, 5629
 Sandrinelli, A., Falomo, R., & Treves, A. 2019, *MNRAS*, **485**, L89
 Stickel, M., Padovani, P., Urry, C. M., Fried, J. W., & Kuehr, H. 1991, *ApJ*, **374**, 431
 Stocke, J. T., Morris, S. L., Gioia, I. M., et al. 1991, *ApJS*, **76**, 813
 Urry, C. M., & Padovani, P. 1995, *PASP*, **107**, 803
 Vermeulen, R. C., Ogle, P. M., Tran, H. D., et al. 1995, *ApJL*, **452**, L5
 Villata, M., & Raiteri, C. M. 1999, *A&A*, **347**, 30
 Villata, M., Raiteri, C. M., Aller, M. F., et al. 2007, *A&A*, **464**, L5
 Villata, M., Raiteri, C. M., Balonek, T. J., et al. 2006, *A&A*, **453**, 817
 Villata, M., Raiteri, C. M., Gurwell, M. A., et al. 2009, *A&A*, **504**, L9
 Villata, M., Raiteri, C. M., Kurtanidze, O. M., et al. 2002, *A&A*, **390**, 407

¹⁵ From the TeVCat catalog: <http://tevcat.uchicago.edu/>.