



Vera C. Rubin Observatory
Data Management

Review of Timeseries Features

Eric C. Bellm

DMTN-118

Latest Revision: 2021-05-28



Abstract

Rubin Observatory will compute timeseries variability features on lightcurves to aid users in identifying objects of interest, both during realtime Alert Production as well as in the annual Data Releases. The Data Products Definition Document (DPDD; LSE-163) allocates space for pre-computed timeseries features, and a sample set is baselined in LDM-151. However, in the subsequent decade the scientific community has made a great deal of further progress in this area. This technote reviews the relevant literature, grouping related features where possible; discusses potential concerns and open questions; and proposes a new baseline feature set.

Change Record

Version	Date	Description	Owner name
1	2021-05-28	Initial release.	Eric Bellm

Document source location: <https://github.com/lstt-dm/dmtn-118>

Contents

1	Current Baseline	1
1.1	Characterization of Periodic Variability	1
1.2	Characterization of Aperiodic Variability	2
2	Motivations for updates to the current baseline	4
3	Other Timeseries Features in the Literature	5
3.1	Characterization of Periodic Variability	5
3.1.1	Multiple Harmonics, Multiple Frequencies	6
3.1.2	Multiple bands	7
3.1.3	Calculations on the Phase-folded lightcurve	7
3.1.4	Nonstationarity	8
3.1.5	Quality Control	8
3.2	Characterization of Aperiodic Variability	9
3.2.1	Amplitudes and Percentiles	9
3.2.2	Other Statistics	9
3.2.3	Characterizing Multi-band Lightcurves	13
3.3	Characterization of Transients	13
3.3.1	Template and Model Fitting	14
4	Open Questions and Practical Concerns	16
4.1	AP and DRP Timeseries features	16
4.2	Multi-band features	17
4.3	Number of data points	17
4.3.1	Period finding	17
4.4	Model Fitting	18
4.5	Solar System Object Features	18
4.6	Interpretability	19
4.7	Storage	19
4.8	Compute features of flux or magnitudes?	19

4.9	Difference or total flux?	19
4.10	Compute AP features using Forced Photometry?	20
4.11	Times	21
4.12	Special Programs	21
4.13	Changing Templates	22
4.14	Evaluation	22
5	Recommendations	23
5.1	Periodic Features	23
5.2	Aperiodic Features	24
6	Proposed Feature Set	24
7	User Interfaces to Timeseries Data	24
A	References	25
B	Acronyms	30

Review of Timeseries Features

1 Current Baseline

Table 2 of the [DPDD](#) allocates space in both of the `DIAObject` and `Object` tables for a total of 6 bands \times (32 periodic + 20 nonperiodic) features = 312 floats. The baseline described in LDM-151 §6.21 is that these features are taken directly from Tables 4 & 5 of Richards et al. (2011). We reproduce the text of this section from [DPDD](#) v4.2 below:

All of these metrics are calculated for both `Objects` ([DPDD](#) table 4, `1cPeriodic` and `1cNonPeriodic`) and `DIAObjects` ([DPDD](#) table 2, `1cPeriodic` and `1cNonPeriodic`). They are calculated and recorded separately in each band. Calculations for `Objects` are performed based on forced point source model fits ([DPDD](#) table 5, `psFlux`). Calculations for `DIAObjects` are performed based on point source model fits to `DIASources` ([DPDD](#) table 1, `psFlux`). In each case, calculation requires the fluxes and errors for all of the sources in the lightcurve to be available in memory simultaneously.

1.1 Characterization of Periodic Variability

- Characterize lightcurve as the sum of a linear term plus sinusoids at three fundamental frequencies plus four harmonics:

$$y(t) = ct + \sum_{i=1}^3 \sum_{j=1}^4 y_i(t|j f_i) \quad (1)$$

$$y_i(t|j f_i) = a_{i,j} \sin(2\pi j f_i t) + b_{i,j} \cos(2\pi j f_i t) + b_{i,j,0} \quad (2)$$

where i sums over fundamentals and j over harmonics.

- Use iterative application of the generalized Lomb-Scargle periodogram, as described in Richards et al. (2011), to establish the fundamental frequencies, f_1 , f_2 , f_3 :
 - Search a configurable (minimum, maximum, step) linear frequency grid with the periodogram, applying a $\log f/f_N$ penalty for frequencies above $f_N = 0.5\langle 1/\Delta T \rangle$, identifying the frequency f_1 with highest power;
 - Fit and subtract that frequency and its harmonics from the lightcurve;
 - Repeat the periodogram search to identify f_2 and f_3 .

- We report a total of 32 floats:
 - The linear coefficient, c (1 float)
 - The values of f_1, f_2, f_3 . (3 floats)
 - The amplitude, $A_{i,j} = \sqrt{a_{i,j}^2 + b_{i,j}^2}$, for each i, j pair. (12 floats)
 - The phase, $PH_{i,j} = \arctan(b_{i,j}, a_{i,j}) - \frac{jf_i}{f_1} \arctan(b_{1,1}, a_{1,1})$, for each i, j pair, setting $PH_{1,1} = 0$. (12 floats)
 - The significance of f_1 vs. the null hypothesis of white noise with no periodic signal. (1 float)
 - The ratio of the significance of each of f_2 and f_3 to the significance of f_1 . (2 floats)
 - The ratio of the variance of the lightcurve before subtraction of the f_1 component to its variance after subtraction. (1 float)

NB the DPDD baselines providing 32 floats, but, since $PH_{1,1}$ is 0 by construction, in practice only 31 need to be stored.

1.2 Characterization of Aperiodic Variability

In addition to the periodic variability described above, we follow Richards et al. (2011) in providing a series of statistics computed from the lightcurve which do not assume periodicity. They define 20 floating point quantities in four groups which we describe here, again with the caveat that future revisions to the DPDD may require changes to this list.

Basic quantities:

- The maximum value of delta-magnitude over delta-time between successive points in the lightcurve.
- The difference between the maximum and minimum magnitudes.
- The median absolute deviation.
- The fraction of measurements falling within 1/10 amplitudes of the median.
- The “slope trend”: the fraction of increasing minus the fraction of decreasing delta-magnitude values between successive pairs of the last 30 points in the lightcurve.

Moment calculations:

- Skewness.

- Small sample kurtosis, i.e.

$$\text{Kurtosis} = \frac{n(n+1)}{(n-1)(n-2)(n-3)} \sum_{i=1}^n \left(\frac{x_i - \bar{x}}{s} \right)^4 - \frac{3(n-1)^2}{(n-2)(n-3)} \quad (3)$$

$$s = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2} \quad (4)$$

- Standard deviation.
- The fraction of magnitudes which lie more than one standard deviation from the weighted mean.
- Welch-Stetson variability index J Stetson (1996), defined as

$$J = \frac{\sum_k \text{sgn}(P_k) \sqrt{|P_k|}}{K},$$

where the sum runs over all K pairs of observations of the object, where sgn returns the sign of its argument, and where

$$P_k = \delta_i \delta_j \quad (5)$$

$$\delta_i = \sqrt{\frac{n}{n-1}} \frac{v_i - \bar{v}}{\sigma_v}, \quad (6)$$

where n is the number of observations of the object, and v_i its flux in observation i . Following the procedure described in Stetson Stetson (1996), the mean is not the simple weighted algebraic mean, but is rather reweighted to account for outliers.

- Welch-Stetson variability index K Stetson (1996), defined as

$$K = \frac{1/n \sum_{i=1}^N N |\delta_i|}{\sqrt{1/n \sum_{i=1}^N N |\delta_i|^2}},$$

where N is the total number of observations of the object and δ_i is defined as above.

Percentiles. Taking, for example, $F_{5,95}$ to be the difference between the 95% and 5% flux values, we report:

- All of $F_{40,60}/F_{5,95}$, $F_{32.5,67.5}/F_{5,95}$, $F_{25,75}/F_{5,95}$, $F_{17.5,82.5}/F_{5,95}$, $F_{10,90}/F_{5,95}$

- The largest absolute departure from the median flux, divided by the median.
- The ratio of $F_{5,95}$ to the median.

QSO similarity metrics, as defined by Butler & Bloom Butler & Bloom (2011):

- χ_{QSO}^2/ν .
- $\chi_{\text{False}}^2/\nu$.

2 Motivations for updates to the current baseline

Both the [DPDD](#) and LDM-151 are clear that the baselined features are provisional and will be updated. Here we briefly review some motivations for doing so now.

First, the Richards et al. (2011) features are now a decade old, and there has been great progress in the field since, with many new classification efforts run on larger and deeper surveys. It is reasonable to review these efforts and distill any new findings. Also, while training machine-learned models on features extracted from the lightcurves remains a popular approach, classification using deep neural nets run directly on lightcurves or their proxies have begun to appear (e.g., Mahabal et al., 2017; Naul et al., 2018; Muthukrishna et al., 2019). What is the appropriate role of LSST-computed features in this rapidly-evolving landscape?

It is also worth noting the ways in which the Richards et al. (2011) study differs from the LSST application. That classification was based on a retrospective analysis of single-band OGLE and *Hipparcos* data, with a particular focus on variable stars. In contrast LSST is a much deeper multi-band survey and will accordingly identify a broader range of classes of astrophysical transients and variables, not all of which will be well-characterized by the baseline feature set. Moreover, features will be calculated and disseminated in alerts in near-real time to aid rapid classification and followup of important events.

Finally we raise the issue of interpretability, which we discuss further in §4.6. The features computed in Richards et al. (2011) were intended primarily as inputs into a machine learning classifier. Because the LSST project will not perform scientific classification, however, its time-series features will be a major way scientists will attempt to identify interesting objects within its data. Additionally, the random forest classifier employed by Richards et al. (2011) readily discards irrelevant features, so there was little incentive to remove redundant or duplicative

features. In contrast the number of LSST features is tightly constrained given the scale of the catalogs.

In sum, these considerations motivate us to look for a set of timeseries features that are orthogonal, interpretable, and span the space of time-domain science LSST data can enable.

3 Other Timeseries Features in the Literature

In this section we review a range of other timeseries features that have been used in the literature. While it is impossible for this list to be exhaustive we have attempted to review the papers for major classification efforts and packages. These included classification of ASAS (Richards et al., 2012b), ASAS-SN (Jayasinghe et al., 2018, 2019a,b, 2020), PS1 (Hernitschek et al., 2016; Villar et al., 2020), HiTS (Martínez-Palomera et al., 2018), VVV (Elorrieta et al., 2016), SNPhotCC (Kessler et al., 2010), the ANTARES broker (Narayan et al., 2018), the PLAsTiCC challenge (Malz et al., 2019; Kessler et al., 2019; Boone, 2019), OGLE (Pashchenko et al., 2018), and ZTF (Sánchez-Sáez et al., 2021; Coughlin et al., 2020). We also reviewed existing packages for featurizing lightcurves, including VARTOOLS (Hartman & Bakos, 2016), FATS (Nun et al., 2015), cesium¹, UPSILoN (Kim & Bailer-Jones, 2016), VaST (Sokolovsky & Lebedev, 2018), astrobase² and tsfresh³.

We attempt to combine related or redundant expressions of the same calculations below. Graham et al. (2017) provides a useful taxonomy of features, suggesting that they may characterize location, scale, variability, morphology, timescales, trends, or agreement with a model. Sokolovsky et al. (2017) and Pashchenko et al. (2018) assess features in the context of determining whether or not a source is variable and note the large degree of covariance among commonly-used metrics.

3.1 Characterization of Periodic Variability

A wide variety of period-finding algorithms have been explored in the literature. These include model-based methods based on the Discrete Fourier Transform (especially the Lomb-Scargle Periodogram and variants; Lomb, 1976; Scargle, 1982; Zechmeister & Kürster, 2009; Vander-

¹<http://cesium-ml.org/>

²<https://astrobase.readthedocs.io>

³<https://tsfresh.readthedocs.io/>

Plas, 2018); phase-folding methods minimizing various scatter measures given a trial period (e.g., Stellingwerf, 1978; Dworesky, 1983; Schwarzenberg-Czerny, 1989; Clarke, 2002; Zucker, 2016); entropy-based methods (Cincotta et al., 1995; Cincotta, 1999; Huijse et al., 2012; Graham et al., 2013a, e.g.); and Bayesian approaches (e.g., Gregory & Loredano, 1992; Wang et al., 2012). Searches for transiting exoplanets typically use the Box Least Squares (BLS) algorithm (Kovács et al., 2002) and variants (e.g., Heller et al., 2019), which fit a narrow transit signal at a variety of trial periods.

In a review, Graham et al. (2013b) report that no single period-finding method is optimal for all classes of periodic object. This suggests that our selection of period finding algorithm(s) should be guided by practical concerns. For LSST, these include the need to provide effective periods and lightcurve characterization for a wide range of astrophysical objects; the need for rapid execution in order to fit within the 60 sec alert latency budget; and the need to fit LSST's temporally sparse multi-band data.

3.1.1 Multiple Harmonics, Multiple Frequencies

Since most variable stars do not have simple sinusoidal lightcurves, authors using Lomb-Scargle period fitting often utilize multi-harmonic and/or multi-frequency models to better fit the lightcurves. The amplitudes and phases of the multiple components help discriminate between classes of variables (e.g., Selam, 2004). We describe a subset below.

Richards et al. (2012b) use a three frequencies and regularized multiple harmonics, using per-lightcurve cross-validation to determine the regularization parameter. Ratios of these frequencies are included as features along with the regularization parameter. Dubath et al. (2011) instead use a hypothesis-testing approach to determine whether additional harmonics are justified⁴. A similar procedure is used by other authors (e.g., Drake et al., 2013; Torrealba et al., 2015) under the name "Adaptive Fourier Decomposition".

Simulation and empirical validation will be necessary to determine the number of frequencies and harmonics which may be meaningfully constrained given the sampling provided by typical LSST cadences in one year.

⁴Although due to boundary-value problems this use of the F-test is uncalibrated: see Protassov et al. (2002).

3.1.2 Multiple bands

Most period fitting conducted in the literature has been performed on survey data taken in one or a small handful of photometric bands. Period fitting accuracy tends to decrease when the number of epochs is smaller than a few tens of data points. LSST will have about 80 epochs yearly divided over six bands for most Wide-Fast-Deep fields; this is too small for effective period fitting in any one band.

Accordingly, several authors have extended the period finding techniques described above to use data in multiple bandpasses simultaneously during period fitting. These include multi-band Lomb Scargle (VanderPlas & Ivezić, 2015), multi-band Analysis of Variance (Mondrik et al., 2015), multi-band Mutual Information (Huijse et al., 2018), as well as hybrid approaches (Saha & Vivas, 2017). Template fitting approaches for specific object types such as RR Lyrae can naturally be conducted in multiple bands (e.g., Stringer et al., 2019).

3.1.3 Calculations on the Phase-folded lightcurve

Richards et al. (2012b) includes features that take the differences in the Lomb-Scargle lightcurve model maxima and minima to detect eclipsing binaries. They also calculate the phase differences between the maxima and minima and the 10th and 90th percentile model slopes, the number of frequencies consistent with 1-day aliases, and whether the ratio of frequencies is consistent with double-mode RR Lyrae. Long et al. (2012) also calculates the same 10th and 90th percentile slopes for the lightcurve after doubling the best-fit period, and adds additional frequency and amplitude ratios.

Dubath et al. (2011) compute several scatter measures in both time and phase space, some that strictly use data values and some that are calculated on residuals from the periodic model fit. These include median differences of consecutive values in both the time- and phase-ordered lightcurves; They also compute the mean squared residuals on the time ordered lightcurve, the phased lightcurve, and the phased lightcurve folded at twice the period. In total there are seven scatter measures with various normalizations. In their classifier, the most valuable of these scatter features is the MAD of the residuals divided by the MAD of the raw light curve.

Kim et al. (2014) calculated von Neumann ratios and range of cumulative sum on the phase-folded lightcurve (§3.2.2).

Jayasinghe et al. (2019b) computes the Lafler-Kinman string-length statistic on both the phased and temporal lightcurves as well as the normalized difference between the two lengths.

3.1.4 Nonstationarity

Some of the most interesting sources LSST finds will be those with periods that change over LSST's ten-year survey. \dot{P} searches are straightforward but too computationally expensive to contemplate at scale. We suggest that identifying candidates to perform such searches on be left to user-generated processing.

Similarly, some sources will show periodic behavior only after detrending long-term variability or by windowing to remove other sources of variation. These specialized analyses are also beyond the scope of routine processing.

3.1.5 Quality Control

Reliable period determination is vital but challenging: uneven sampling, sensitivity to the choice of frequency grid, outliers, non-stationary processes, and frequency-dependent noise properties can all yield spurious periodicities.

It is desirable to have a measure of the probability that the computed period is a false alarm: that is, how likely is it that a non-periodic source would yield the observed periodogram peak? While the Lomb-Scargle periodogram provides an analytic False Alarm Probability (FAP) of a peak at a *given* frequency (Scargle, 1982) this is not the relevant FAP of a finding a spurious peak at *any* frequency (VanderPlas, 2018, and references therein). Improved analytic estimates for the FAP over all frequencies under certain limits (Baluev, 2008), but better constraints typically require computationally expensive bootstrap resampling (e.g., Ivezić et al., 2014), although use of extreme-value statistics can reduce the number of required bootstrap evaluations (Süveges, 2012). Recent work (Delisle et al., 2020) extends analytic FAP estimates to handle the astrophysically-common case of temporally-correlated variability.

However, the FAP does not tell us whether the identified period is correct. Often the highest periodogram peak is a fraction or multiple of the true period, particularly for eclipsing binaries with double-peaked lightcurves. Dubath et al. (2011) provides several empirical measures to assess period quality, including the largest phase gap in the folded lightcurve and MAD and

percentile values of the scatter of the residuals about the lightcurve phased at the Lomb-Scargle period as well as twice its value. These are also adopted by (Richards et al., 2012b). Sánchez-Sáez et al. (2021) computes a “periodogram pseudo-entropy” from peaks in the periodogram as well as power ratios of the selected period and ratios thereof.

Given that no single period-finding method is optimal for all classes of sources (Graham et al., 2013b), it would be valuable to report periods derived from multiple period-finding algorithms. However, this may not be computationally feasible, particularly within the required 60 sec alert production latency.

3.2 Characterization of Aperiodic Variability

3.2.1 Amplitudes and Percentiles

A wide variety of combinations and scalings of percentile amplitudes are reported in the literature. For example, Kinemuchi et al. (2006) uses a magnitude ratio of (max–median) / (max–min)), sometimes known as the “M-test statistic.” Kim & Bailer-Jones (2016) calculates the difference of the 75% and 25% values (the inter-quartile range).

We suggest it may be more expedient to simply report the raw flux percentiles per band and let users construct differences and ratios as desired, rather than using the complex set of normalized difference ratios defined in Richards et al. (2011) and LDM-151.

3.2.2 Other Statistics

Some commonly-computed statistics include reduced χ^2 for a constant model:

$$\chi^2/\text{d.o.f.} = \frac{1}{N-1} \sum_{i=1}^N \left(\frac{m_i - \bar{m}}{\sigma_i} \right)^2 \quad (7)$$

This may be used to derive the probability that the variability is unlikely to be due to chance variations from a constant source (e.g., McLaughlin et al., 1996).

Enoch et al. (2003) proposed a Robust Median Statistic:

$$\text{RoMS} = (N - 1)^{-1} \sum_{i=1}^N \frac{|m_i - \text{median}(m_i)|}{\sigma_i} \quad (8)$$

Excess variance σ_{sys} is sometimes used to capture the amount of variability not represented by statistical error bars (e.g., in Bernstein et al., 2018):

$$\begin{aligned} \sigma_{\text{sys}}^2 &\equiv \left\langle \Delta m_i^2 - \sigma_{\text{stat},i}^2 \right\rangle \\ \Delta m_i &\equiv \frac{m_i - \bar{m}}{\sqrt{1 - w_i / \sum w_j}} \\ w_i &\equiv \sigma_{\text{stat},i}^{-2} \\ \bar{m} &\equiv \frac{\sum w_j m_j}{\sum w_j} \end{aligned} \quad (9)$$

The normalized excess variance can be calculated (e.g., Nandra et al., 1997; Allevato et al., 2013) as:

$$\sigma_{\text{sys,norm}}^2 \equiv \frac{\sigma_{\text{sys}}^2}{N \bar{m}^2} \quad (10)$$

Richards et al. (2012b) compute a robust skew measure g_{skew} designed for objects that rapidly decrease in brightness:

$$(\text{med}(m) - \text{med}(m[0 : p])) + (\text{med}(m) - \text{med}(m[p : 10]))$$

where $p = 0.03$.

Kim et al. (2011) calculate the number of objects above and below a variability threshold computed from the autocorrelation. They also calculate the “range of a cumulative sum” defined as

$$\begin{aligned} R_{\text{CS}} &= \max(S) - \min(S) \\ S_l &= \frac{1}{N\sigma} \sum_{i=1}^l (m_i - \bar{m}). \end{aligned} \quad (11)$$

They also compute the ratio of the standard deviation to the mean magnitude. They adapt a

Con metric from Wozniak (2000) that assesses how many sets of three consecutive data points are more than 2σ outliers in the same direction. Finally they use the von Neumann ratio

$$\eta = \frac{1}{(N-1)\sigma^2} \sum_{i=1}^{N-1} (m_{i+1} - m_i)^2 \quad (12)$$

which assesses the smoothness of a lightcurve by calculating the squared differences of consecutive points.

Kim et al. (2014) proposed an alternative von Neumann ratio to account for unequal sampling:

$$\eta^e = \bar{\omega} (t_{N-1} - t_1)^2 \frac{\sum_{i=1}^{N-1} w_i (m_{i+1} - m_i)^2}{\sigma^2 \sum_{i=1}^{N-1} w_i} \quad (13)$$

$$w_i = \frac{1}{(t_{i+1} - t_i)^2}.$$

Mowlavi (2014) discusses the Abbe value, $\mathcal{A} = \eta/2$, and performs a sliding window computation of it on sub-regions of the lightcurve. Comparing the mean of these subsample Abbe values \mathcal{A}_{sub} to that of the whole timeseries gives the “excess Abbe value”:

$$\mathcal{E}_{\mathcal{A}} \equiv \overline{\mathcal{A}_{\text{sub}}} - \mathcal{A} \quad (14)$$

which can distinguish transient, pulsating, trended, and featureless lightcurves in the $\mathcal{E}_{\mathcal{A}}$ vs. \mathcal{A} plane for 100-day window timescales.

Kim & Bailer-Jones (2016) computes the ratio of the squared sum of residuals above and below the median, which Jayasinghe et al. (2019b) reports is effective in identifying eclipsing binaries. They also include the Shapiro-Wilk statistic, which tests for normality.

Tamuz et al. (2006) defines an “alarm” statistic for detecting runs on residuals $r_{i,j}$:

$$\mathcal{A}_{\text{alarm}} = \frac{1}{\chi^2} \sum_{i=1}^M \left(\frac{r_{i,1}}{\sigma_{i,1}} + \frac{r_{i,2}}{\sigma_{i,2}} + \dots + \frac{r_{i,k}}{\sigma_{i,k_i}} \right)^2 - \left(1 + \frac{4}{\pi} \right) \quad (15)$$

where i denotes the i th “run” of consecutive residuals of the same sign and j is the j th measurement in that run. This statistic is implemented in VARTOOLS for a constant model. Sokolovsky

et al. (2017) describes some related variants, including the S_B variability statistic of Arellano Ferro et al. (2012).

Stetson (1996) developed an L index which is a weighted product of the J and K indices already included in our baseline feature set. Sokolovsky et al. (2017) describes variants of the Welch-Stetson I Welch & Stetson (1993) and Stetson J , K , and L indices which relax assumptions that the observations are taken in simultaneous pairs.

tsfresh includes a range of statistics for generic non-astronomical timeseries, including approximate entropy⁵, which quantifies irregularity in timeseries data; various tests and measures of autocorrelation and nonstationarity; and number of mean crossings.

Mislis et al. (2016) presents two features later adopted by Narayan et al. (2018), an integral of the autocorrelation over all lags τ :

$$A_1 = \left| \sum_{r=1}^n \left(\frac{1}{(n-\tau) \cdot \text{rms}^2} \sum_{i=1}^{n-r} (x_i - \bar{x}) (x_{i+\tau} - \bar{x}) \right) \right| \quad (16)$$

and a modification of the Shannon Entropy comparing the observed CDF of the data to a normal distribution.

Sánchez-Sáez et al. (2021) computes a “Mexican Hat Power Spectrum” to assess the lightcurve variance at 10 and 100 day timescales.

Many classes of accreting systems—especially AGN but also cataclysmic variables, X-ray binaries, and Young Stellar Objects—show stochastic variability.

(Kozłowski, 2016, and references therein) describe the use of the Structure Function to capture the variability of the data as a function of lag time δt . This distribution may be characterized by a power-law index and a decorrelation timescale at which the Structure Function slope flattens into white noise. The authors caution about conflicting definitions and parameterizations of the Structure Function in the literature. (Hernitschek et al., 2016) uses Gaussian Processes to derive a multi-band Structure Function, while Hu & Tak (2020) uses a Kalman filtering approach to determine the likelihood of a multi-band Damped Random Walk model with $O(k^3 n)$

⁵https://en.wikipedia.org/wiki/Approximate_entropy

complexity. (Graham et al., 2014) proposes an alternative metric, the Slepian Wavelet Variance. A general family of Continuous Autoregressive Moving Average (CARMA) models can also be used to model the stochastic variability (Moreno et al., 2019, and references therein). Sánchez-Sáez et al. (2021) and De Cicco et al. (2021) included Damped Random Walk, Structure Function, and Irregular Autoregressive (CAR(1)) fits in their classification of ZTF alerts and VLT data, respectively.

Several distinct classes of astronomical sources show outbursts (or dips) relative to a quiescent magnitude; these include cataclysmic variables and other compact binaries, stellar flares, Young Stellar Objects, and AGN. The *Lasair* broker uses exponential moving averages to detect outbursts (Roy Williams, priv. comm.):

$$EMA(t) = f * EMA(t_p) + (1 - f) * m(t) \quad (17)$$

$$f = \exp(-(t - t_p)/\tau)$$

where t_p is the previous detection time and τ is a timescale. *Lasair* uses τ values of 2 days, 8 days, and 28 days. Other approaches for detecting outbursts includes Bayesian Blocks (Scargle et al., 2013) and various prescriptions for detecting consecutive “runs” (e.g., Chang et al., 2015).

3.2.3 Characterizing Multi-band Lightcurves

Kim et al. (2014) calculates quantiles and the von Neumann ratio on a $B - R$ color lightcurve derived from a survey with near-simultaneous multi-band measurements.

3.3 Characterization of Transients

While the boundary between transient events and aperiodic variability can be fuzzy, the baselined features from Richards et al. (2011) were derived solely for applications to variable stars. Rapid identification and classification of transients in the LSST alert stream will thus benefit from their own specialized features. We review here several transient classification approaches from the literature⁶.

⁶From <https://www.kaggle.com/michaelapers/the-plasticc-astronomy-classification-demo> and <https://github.com/villrv/TransientClassifierTable>

The redshift of the transient host galaxy is a key input for transient classification. LSST will not provide directly a single redshift for a given transient. Rather, each `DIAObject` will include the IDs of the three nearest star-like and galaxy-like `Objects` from the most recent LSST Data Release as well as the three nearest extended objects from a low-redshift galaxy catalog (`DPDD`). Science users or brokers can use these ids to obtain photometric redshifts from LSST Data Releases if they have Data Rights. An additional crossmatch to a low-redshift galaxy catalog is under study. DMTN-049 discusses photometric redshifts in more detail.

3.3.1 Template and Model Fitting

Lightcurve fitting is common in transient classification. Some authors use software (e.g., Jha et al., 2007; Kessler et al., 2009; Sako et al., 2011; Barbary, 2014; Guillochon et al., 2018) to fit templates or models derived from observed transients (e.g., Nugent et al., 2002; Guy et al., 2007; Conley et al., 2008; Kessler et al., 2019; Vincenzi et al., 2019).

Others use analytic models meant to capture typical transient lightcurve shapes.

Newling et al. (2011) fit a parametric function with free parameters of explosion time, amplitude, relative rise and decay times, temporal stretch, and a tail decay function $\Psi(t)$:

$$F(t) = A \left(\frac{t - \phi}{\sigma} \right)^k \exp \left(-\frac{t - \phi}{\sigma} \right) k^{-k} e^k + \Psi(t). \quad (18)$$

Bazin et al. (2011) developed a parameterization with the goal of identifying SNe Ia:

$$f^k(t) = A^k \frac{e^{-\left(t-t_0^k\right)/\tau_{fall}^k}}{1 + e^{-\left(t-t_0^k\right)/\tau_{rise}^k}} + c^k. \quad (19)$$

Karpenka et al. (2013) developed a generic parameterization that can fit double-peaked lightcurves:

$$f(t) = A \left[1 + B (t - t_1)^2 \right] \frac{e^{-(t-t_0)/T_{fall}}}{1 + e^{-(t-t_0)/T_{rise}}}. \quad (20)$$

Villar et al. (2019) developed an analytic model for generic explosive transient lightcurves, including a plateau phase (but that does not capture the second peak in Type Ia SNe directly):

$$F = \begin{cases} \frac{A + \beta(t - t_0)}{1 + e^{-(t - t_0)/\tau_{\text{rise}}}} & t < t_1 \\ \frac{(A + \beta(t_1 - t_0))e^{-(t - t_1)/\tau_{\text{fall}}}}{1 + e^{-(t_0 - t)/\tau_{\text{rise}}}} & t \geq t_1 \end{cases} \quad (21)$$

Sánchez-Sáez et al. (2021) modified the parametric model of Villar et al. (2019) to smoothly transition between the two piecewise functions:

$$F = \frac{A \left(1 - \beta' \frac{t - t_0}{t_1 - t_0}\right)}{1 + \exp\left(-\frac{t - t_0}{\tau_{\text{rise}}}\right)} \cdot \left[1 - \sigma\left(\frac{t - t_1}{3}\right)\right] + \frac{A(1 - \beta') \exp\left(-\frac{t - t_1}{\tau_{\text{fall}}}\right)}{1 + \exp\left(-\frac{t - t_0}{\tau_{\text{rise}}}\right)} \cdot \left[\sigma\left(\frac{t - t_1}{3}\right)\right] \quad (22)$$

where $\sigma(t)$ is the sigmoid function.

Microlensing events have characteristic shapes which can be fit using standard optimization algorithms. VARTOOLS implements a basic microlensing model from Wozniak et al. (2001).

Finally, some authors employ non-parametric models. Richards et al. (2012a) and Ishida & de Souza (2013) used a nonparametric cubic spline. Varughese et al. (2015) and Lochner et al. (2016) decomposed the lightcurve using wavelets after mapping to a uniform temporal grid using cubic splines and Gaussian Processes, respectively. Revsbech et al. (2018) used Gaussian Processes for data augmentation and Boone (2019) introduced a Gaussian Process model for arbitrary transients that modeled both the temporal and spectral dimensions, enabling simultaneous multi-band fitting. For our purposes, the difficulty with nonparametric representation is that they are difficult to represent compactly, although various features may be extracted from the nonparametric model (e.g., Boone, 2019).

For explosive transients, key generic observables to capture include the slope of the rising and falling lightcurve, estimates of the source color, and estimates of the time of the peak. However, these may vary from band to band. In Wide-Fast-Deep the sparse sampling in indi-

vidual filters may not provide enough points to perform fits in individual bands. Additionally, some form of interpolation may be required to estimate source color as a function of time depending on the distribution of filters within the adopted LSST cadence.

Information about non-detections can also be important in understanding a transient's evolution and distinguishing it from other classes of source. Sánchez-Sáez et al. (2021) included several nondetection features in their ZTF alert classifier.

Villar et al. (2020) trained an autoencoder to compress Gaussian-Process-interpolated supernova lightcurves into a low-dimensional latent space. These features were then combined with per-band times to rise to peak, times to fall from peak, peak absolute magnitudes, median post-peak slope, and light-curve integral (each calculated from the interpolated GP lightcurves) to train a random-forest classifier.

4 Open Questions and Practical Concerns

In this section we highlight open issues that warrant further study.

4.1 AP and DRP Timeseries features

Should AP's DIAobject features be the same as the DRP DIAobject and Object features, as assumed in the current baseline? Without AP's latency requirements, more computationally expensive features (or simply more features) could be computed. DRP will compute features for the full LSST lightcurve, so features sensitive to more subtle signals could be valuable.

AP timeseries features will not need to identify variable objects (e.g., Pashchenko et al., 2018), as in AP only objects that vary relative to the template image will produce DIASources. In DRP, however, we will compute timeseries features on all DIAObjects and Objects regardless of their variability.

Having disparate feature sets between AP and DRP could cause some user confusion, but our judgment is that this is a minor issue. We suggest that a larger set of features be contemplated for DRP (subject to sizing model considerations), but that they be a superset of those used for AP so that users can compare AP and DRP features that are common.

4.2 Multi-band features

Most timeseries features to date are computed on single-band timeseries data. The current baseline calls for LSST to compute single-band features independently on each LSST band. *Are there new features we could develop that take advantage of LSST's sparse multi-band data?* Multi-band period fitting has already been developed. Model-dependent lightcurve fits could provide interpolated color estimates for transients. Are there other timeseries features that could be usefully generalized from their current single-band form?

One option would be to perform single-band fitting when there are enough data points to warrant it, and to fall back to multi-band fitting in other cases. The disadvantage to this approach is that it is wasteful from a storage perspective (the same fit is duplicated over six per-band columns).

4.3 Number of data points

When a new DIAobject is identified in AP, it will only have one DIASource. *Are there threshold numbers of epochs (per band) we should require before computing some timeseries features?* Even basic statistics like the mean require 2–3 points to provide information, and more complex features like period finding need tens of epochs. Since many false positives will only have a few detections, omitting them from feature computation would provide savings.

We propose to empirically determine the minimum numbers of DIASources required to obtain meaningful results for various features. Features not computed will have values set to NULL.

4.3.1 Period finding

The LSST main survey has 825 fiducial visits per pointing, divided among six bands. In AP, the number of epochs available for period fitting in a single band using the 12 month DIASource history (10–30) is at the very lower limit of what is needed for effective period fitting. This suggests that a multi-band period fit will be required for AP.

DRP will fit timeseries features using the entire DIASource history, so after several years of the survey single-band period finding could be effective. Are there astrophysical sources where we expect the period to differ by band, or where the assumptions of multi-band period fitting

will break down?

DRP will see some sources where the period may change over the lifetime of LSST. Are there enough sources where this is a problem to justify limiting the temporal range of data input to period finding?

4.4 Model Fitting

Fits to theoretical models for a wide variety of transients, variables, and moving objects can provide great insight into their astrophysics. Models of supernovae, other transients, microlensing events, eclipsing binaries, spotted stars, stellar flares, and more are widely applied in the literature.

Should LSST perform science-specific model fits?

We argue that in general these fits are better done by science users on self-selected subsamples of the LSST data rather than as part of generic feature computation on all `DIAObjects` and `Objects`. First, the very nature of the models is that they are only meaningful for a small subset of the total number of objects LSST has data for; even identifying the relevant subset requires classification which is beyond LSST scope. Additionally, model fitting is typically computationally expensive, and is unlikely to be compatible with the 60-second alert latency budget. Moreover, it is not clear there is sufficient space to store all of the resulting best-fit model parameters, their uncertainties, and the goodness of fit. Finally, the choice of models to fit, their parameter constraints, and so on are best left to specialists.

Accordingly, we recommend features that characterize broad classes of generic events: e.g., the rise time of a transient calculated from a nonparametric fit rather than a SALT2 SN Ia template.

4.5 Solar System Object Features

The schema proposed in RFC 620 for the `SSObject` catalog allocates the same amount of space for periodic features as for `DIAObjects`. From a code reuse perspective, having the same set of features for both `DIAObjects` and `SSObjects` is simple. However, it is not obvious that features developed for transients and AGN are useful to compute on `SSObjects`.

Additionally, we note that for the features to be useful for SSObjects they should be computed on phase-corrected SSSource data points rather than the original DIASources.

4.6 Interpretability

All of the timeseries features computed will be available for user query. This suggests that we prefer features that are more intelligible to humans where possible, in order to aid in construction of appropriate queries and filters. We expect that the highest-performance machine-learning classifiers will work directly on the data products themselves rather than on the precomputed features, so feature interpretability is more valuable than performance in a hypothetical classifier.

4.7 Storage

We have budgeted for floats, but most statistics don't have or need many digits of floating point precision. We could cast some features as integers either to save space or to make room for additional features. However, this conversion is lossy and imposes an additional usability burden on users, who must remember when and how to convert any fields to floating point⁷. Additionally the range of the feature must be well-understood to prevent overflows.

4.8 Compute features of flux or magnitudes?

Most optical classification programs in the literature have featurized lightcurves reported in magnitudes. Given the complications magnitudes impose with negative detections (which will be common in the forced photometry table, for instance) we will compute our features in flux space. It is worth noting that for some features this may lead to altered behavior compared to the literature; we should ensure that the statistical sense of the features we compute is preserved.

4.9 Difference or total flux?

Each DIASource has multiple flux measurements associated with it. The most important values are psFlux, the point-source flux measured on the difference image, and totFlux, a *forced*

⁷or use a provided convenience function

PSF flux measured at the DIASource centroid on the calibrated (unsubtracted) visit image. Different science cases may require that distinct timeseries features be calculated on lightcurves built from different fluxes. Features that try to capture the rise time of a supernovae will obviously use a `psFlux` lightcurve, as there is no relevant flux in the template. For a periodogram of a variable star, however, the total flux including the flux in the reference may be more appropriate.

It may be suggested that computing timeseries features on fluxes rather than on magnitudes makes this distinction irrelevant. However, since the templates used for subtraction are updated in every data release, this means that the range of `psFlux` values for variable stars will change with the images used in the templates. This will create jumps when AP processing switches to new templates after each data release. These offsets are different for each variable object because their phases and flux levels are uncorrelated. Section 4.13 discusses some possible mitigations.

It may be preferable to compute variable star features on the forced photometry measurement (`totFlux`) on the Processed Visit Image, which would provide the correct absolute flux values no matter the template⁸. However, these measurements will be less reliable in crowded fields, and they are susceptible to centroid errors in the triggering DIASource.

4.10 Compute AP features using Forced Photometry?

LSST obtains several kinds of forced photometry measurements during Alert Production. When a new `DIAObject` is created, “precovery” forced photometry is obtained for any images at that position observed in the past 30 days. And each time a field is observed, LSST will compute forced photometry at the positions of all overlapping `DIAObjects` that have been detected in the last 12 months. In both cases forced photometry is performed both on the difference image and the unsubtracted processed visit image, so the discussion in §4.9 also applies to forced photometry.

Should we use forced photometry measurements (either on difference images or PVI) to compute timeseries features in AP? If we only compute features when DIASources are detected, variable star photometry will be highly biased, as no DIASources will be present when the variable star

⁸The Alerce ZTF alert classifier described in Sánchez-Sáez et al. (2021) uses a total flux estimate for its variable star features. As ZTF forced photometry measurements are not available, they use the sum of the difference flux and the template flux of the spatially coincident template source, if such a source exists.

flux is near the template value. However, PSF forced photometry fluxes are sensitive to centroiding errors, whether due to bad image subtractions or errors in modeling proper motion. Moreover, for objects that have faded significantly many forced photometry measurements will simply record zero flux with background noise, which could dilute the signal for some features. Sánchez-Sáez et al. (2021) achieved reasonable performance in classifying variable stars from ZTF alerts despite the absence of forced photometry measurements.

There is also a subtlety in interpreting the alert contents. Alerts are only triggered by `DIA-Source` *etections*, but if we use forced photometry lightcurves for timeseries features the latest measurement will not be identical to the triggering `DIASource`. Also, `DIAObjects` would require updating each time a new forced photometry measurement was made, which is not in the current processing plan.

If we did decide that we wanted to update `DIAObject` features with forced photometry results (and no triggering `DIASource`) we could do so in a more leisurely way—these would be subject to `L1PublicT` rather than `OTT1` requirements.

In `DRP` uniform forced photometry lightcurves will be available for all `Objects` and `DIAObjects`, so it will make sense to use these to compute the timeseries features.

4.11 Times

Period-finding and other computations for short-timescale events can be biased by light-travel time across the solar system. Accordingly our feature computation should use barycentered times. Pipelines will natively report TAI times. *Should we also store the barycentered times in the `DIASource` record?* The transformation is straightforward to compute; however, if the barycentered time is omitted users will have to recognize that they cannot reproduce our features without first computing the barycentered times.

4.12 Special Programs

Observations of Special Programs might benefit from different timeseries features to capitalize on the potential changes in image characteristics, time sampling, and total number of observations. *Would special programs benefit from unique timeseries features, and could this be accommodated technically?* As discussed in `DMTN-065`, Special Programs data will be pro-

cessed using the same algorithms used for data from the main survey, but the algorithms can be reconfigured. In practice this might mean adding additional frequencies to a Lomb-Scargle periodogram computation, for example, but not adding entirely new classes of model fits.

4.13 Changing Templates

AP expects to change templates after each major data release, both to improve the quality of the templates and to minimize false detections due to objects with larger proper motion. Template changes will introduce offsets in the difference fluxes for most objects. As we expect Data Releases will occur roughly annually, and we compute lightcurve histories in Alert Production on twelve months of data, this means that there will always be a template offset present in our feature computation.

There are a range of possible solutions to this problem. We could determine offsets for all `DIAObject`s by processing the same science exposures with the new and old templates and measuring the change in the `DIASource` flux, but it is unclear how reliable this procedure would be in practice. A more robust possibility for deriving the offset would be to conduct forced photometry at the `DIAObject` locations in the old and new templates and difference their fluxes, or to difference the old and new templates and perform forced photometry at the `DIAObject` locations on the difference of the templates. In either case it would be necessary to apply these offsets consistently and transparently during feature computation for twelve months, before they would be superseded by new values after the next data release.

4.14 Evaluation

Given the options and questions above, *how can we choose which features to implement?*

Superficially, a large-scale data challenge seems like an appealing way to find the “best” features. However, as discussed in §4.6, raw performance by machine-learning classifiers is not the only goal for these features, since they are also used by humans to write queries and alert filters. Additionally, there is no truly LSST-like dataset in existence today to conduct the data challenge on, and the most realistic simulation effort to date (PLaSTICC) only simulated a subset of the wide-ranging event classes LSST will see. Finally, there is no effort available in the LSST construction project to support such an effort, and it is unclear that community members are prepared to undertake such a large project without substantial financial support.

Accordingly we suggest using this document as a starting point and a repository for discussion with the LSST Science Collaborations, and will include the evolving proposed feature set below as an appendix. We suggest that community-developed simulations and precursor survey data be used to scientifically validate the features proposed here and their implementation, instead of using them in a data challenge format to actually *select* the features.

Concretely, we suggest that the Project personnel, in coordination with members of the Science Collaborations, use a limited set of *fiducial* object classes, simulated on the baseline cadence, to tune the proposed feature set (e.g., determining the number of period frequencies and harmonics that may be successfully fit; identifying threshold numbers of epochs needed for some features). Broad community validation of the features would occur during the Data Previews (although the temporal baseline and cadences will not match that expected during the full survey).

A relevant and open question is, *Will LSST compute the same features over the ten year survey?* Maintaining the same feature set provides continuity and simplifies the continuously-updated PPDB. However, it requires us to correctly develop the LSST feature set before having seen any LSST data, which seems difficult. One option would be to only use a portion of the allocated space in the early survey, so that additional features could be added in later years of operations without removing features previously calculated.

5 Recommendations

Here we summarize the questions listed above and our recommendations.

5.1 Periodic Features

Since LSST images are spread among six filters, we expect that best results will be achieved by using a multi-band period search technique. We recommend providing an Adaptive Fourier Decomposition rather than a nonparameteric period estimate, as the Fourier parameters may be used to identify variable types.

We will include an analytic estimate of the False Alarm Probability (e.g., Baluev, 2008; Delisle et al., 2020) and goodness of fit measures.

5.2 Aperiodic Features

For flexibility, we recommend storing raw percentiles rather than combinations thereof—if desired, user-defined functions may be provided to simplify computations (e.g., of the 5%–95% amplitude).

We recommend use of robust statistics where possible due to the unavoidable presence of outliers. We will investigate the tradeoffs in identifying and potentially excluding outliers from some statistics (e.g., Pashchenko et al., 2018).

To provide useful features for the widest range of transient and variable objects, we should include measures of stochastic variability and transient characterization. Given the wide scientific community which uses these features, generic characterization of transients will be more broadly useful than detailed models specializes for individual transient types.

6 Proposed Feature Set

To be determined after consultation with the science community.

A draft proposal for features can be found at <https://ls.st/fkr>. It will be incorporated into this technote after an initial phase of community feedback.

7 User Interfaces to Timeseries Data

Users will access LSST timeseries data via ADQL queries through the Science Platform’s VO APIs. Concretely, they will receive `VOTables` of `DIAObjects`, `DIASources`, etc. Depending on the parameters of the query (e.g., a large spatial cone search), multiple distinct astrophysical objects may be included in the response. It will be the user’s responsibility to handle multiple distinct returned objects appropriately.

In Python environments such as the Science Platform’s Notebook aspect, the `pyVO` package can be used to transform the returned `VOTable` into `astropy Tables`.

The Rubin Observatory pipelines that compute timeseries features do not do so on `VOTables`

or astropy Tables. Accordingly, to re-run the timeseries feature computation code on data returned by user queries it will be necessary to translate the results into the appropriate pipeline-internal representations. We expect to provide basic user tutorials for converting query payloads back into these internal data structures. Thus we expect users to be able to functionally reproduce our feature computations by running the feature code in the Science Platform on data returned from TAP queries.

We do not plan at this time to develop or deploy a “Timeseries object” abstraction. We note that similar efforts are underway in astropy and the VO, and we are continuing to monitor those developments.

A References

References

- Allevato, V., Paolillo, M., Papadakis, I., Pinto, C., 2013, *ApJ*, 771, 9 (arXiv:1202.3786), doi:10.1088/0004-637X/771/1/9, ADS Link
- Arellano Ferro, A., Bramich, D.M., Figuera Jaimes, R., Giridhar, S., Kuppuswamy, K., 2012, *MNRAS*, 420, 1333 (arXiv:1111.0216), doi:10.1111/j.1365-2966.2011.20119.x, ADS Link
- Baluev, R.V., 2008, *MNRAS*, 385, 1279
- Barbary, K., 2014, URL <https://doi.org/10.5281/zenodo.592747>, doi:10.5281/zenodo.592747
- Bazin, G., Ruhlmann-Kleider, V., Palanque-Delabrouille, N., et al., 2011, *A&A*, 534, A43 (arXiv:1109.0948), doi:10.1051/0004-6361/201116898, ADS Link
- Bernstein, G.M., Abbott, T.M.C., Armstrong, R., et al., 2018, *PASP*, 130, 054501 (arXiv:1710.10943), doi:10.1088/1538-3873/aaa753, ADS Link
- Boone, K., 2019, *AJ*, 158, 257
- Butler, N.R., Bloom, J.S., 2011, *AJ*, 141, 93 (arXiv:1008.3143), doi:10.1088/0004-6256/141/3/93, ADS Link
- Chang, S.W., Byun, Y.I., Hartman, J.D., 2015, *ApJ*, 814, 35 (arXiv:1510.01005), doi:10.1088/0004-637X/814/1/35, ADS Link

- Cincotta, P.M., 1999, MNRAS, 307, 941
- Cincotta, P.M., Mendez, M., Nunez, J.A., 1995, ApJ, 449, 231
- Clarke, D., 2002, A&A, 386, 763
- Conley, A., Sullivan, M., Hsiao, E.Y., et al., 2008, ApJ, 681, 482 (arXiv:0803.3441), doi:10.1086/588518, ADS Link
- Coughlin, M.W., Burdge, K., Duev, D.A., et al., 2020, arXiv e-prints, arXiv:2009.14071
- De Cicco, D., Bauer, F.E., Paolillo, M., et al., 2021, A&A, 645, A103
- Delisle, J.B., Hara, N., Ségransan, D., 2020, A&A, 635, A83
- Drake, A.J., Catelan, M., Djorgovski, S.G., et al., 2013, ApJ, 763, 32
- Dubath, P., Rimoldini, L., Süveges, M., et al., 2011, MNRAS, 414, 2602
- Dworetzky, M.M., 1983, MNRAS, 203, 917
- Elorrieta, F., Eyheramendy, S., Jordán, A., et al., 2016, A&A, 595, A82
- Enoch, M.L., Brown, M.E., Burgasser, A.J., 2003, AJ, 126, 1006
- Graham, M., Drake, A., Djorgovski, S.G., Mahabal, A., Donalek, C., 2017, In: European Physical Journal Web of Conferences, vol. 152, 03001
- [DMTN-065]**, Graham, M., Jurić, M., Lim, K.T., Bellm, E., 2018, *Data Management and LSST Special Programs*, DMTN-065, URL <https://dmtn-065.lsst.io>, LSST Data Management Technical Note
- Graham, M.J., Drake, A.J., Djorgovski, S.G., Mahabal, A.A., Donalek, C., 2013a, MNRAS, 434, 2629
- Graham, M.J., Drake, A.J., Djorgovski, S.G., et al., 2013b, MNRAS, 434, 3423
- Graham, M.J., Djorgovski, S.G., Drake, A.J., et al., 2014, MNRAS, 439, 703 (arXiv:1401.1785), doi:10.1093/mnras/stt2499, ADS Link
- [DMTN-049]**, Graham, M.L., 2018, *LSST DRP (Level 2) Catalog Photometric Redshifts*, DMTN-049, URL <https://dmtn-049.lsst.io>, LSST Data Management Technical Note
- Gregory, P.C., Loredó, T.J., 1992, ApJ, 398, 146

- Guillochon, J., Nicholl, M., Villar, V.A., et al., 2018, *ApJS*, 236, 6
- Guy, J., Astier, P., Baumont, S., et al., 2007, *A&A*, 466, 11 (arXiv:astro-ph/0701828), doi:10.1051/0004-6361:20066930, ADS Link
- Hartman, J.D., Bakos, G.Á., 2016, *Astronomy and Computing*, 17, 1
- Heller, R., Rodenbeck, K., Hippke, M., 2019, *A&A*, 625, A31
- Hernitschek, N., Schlafly, E.F., Sesar, B., et al., 2016, *ApJ*, 817, 73
- Hu, Z., Tak, H., 2020, *AJ*, 160, 265
- Huijse, P., Estevez, P.A., Protopapas, P., Zegers, P., Principe, J.C., 2012, *IEEE Transactions on Signal Processing*, 60, 5135
- Huijse, P., Estévez, P.A., Förster, F., et al., 2018, *ApJS*, 236, 12
- Ishida, E.E.O., de Souza, R.S., 2013, *MNRAS*, 430, 509 (arXiv:1201.6676), doi:10.1093/mnras/sts650, ADS Link
- Ivezić, Ž., Connelly, A.J., VanderPlas, J.T., Gray, A., 2014, *Statistics, Data Mining, and Machine Learning in Astronomy*, ADS Link
- Jayasinghe, T., Kochanek, C.S., Stanek, K.Z., et al., 2018, *MNRAS*, 477, 3145
- Jayasinghe, T., Stanek, K.Z., Kochanek, C.S., et al., 2019a, *MNRAS*, 486, 1907
- Jayasinghe, T., Stanek, K.Z., Kochanek, C.S., et al., 2019b, *MNRAS*, 485, 961
- Jayasinghe, T., Stanek, K.Z., Kochanek, C.S., et al., 2020, *MNRAS*, 491, 13
- Jha, S., Riess, A.G., Kirshner, R.P., 2007, *ApJ*, 659, 122 (arXiv:astro-ph/0612666), doi:10.1086/512054, ADS Link
- [LSE-163]**, Jurić, M., et al., 2017, *LSST Data Products Definition Document*, LSE-163, URL <https://ls.st/LSE-163>
- Karpenka, N.V., Feroz, F., Hobson, M.P., 2013, *MNRAS*, 429, 1278 (arXiv:1208.1264), doi:10.1093/mnras/sts412, ADS Link
- Kessler, R., Bernstein, J.P., Cinabro, D., et al., 2009, *PASP*, 121, 1028
- Kessler, R., Bassett, B., Belov, P., et al., 2010, *PASP*, 122, 1415 (arXiv:1008.1024), doi:10.1086/657607, ADS Link

- Kessler, R., Narayan, G., Avelino, A., et al., 2019, *PASP*, 131, 094501 (arXiv:1903.11756), doi:10.1088/1538-3873/ab26f1, ADS Link
- Kim, D.W., Bailer-Jones, C.A.L., 2016, *A&A*, 587, A18
- Kim, D.W., Protopapas, P., Byun, Y.I., et al., 2011, *ApJ*, 735, 68
- Kim, D.W., Protopapas, P., Bailer-Jones, C.A.L., et al., 2014, *A&A*, 566, A43
- Kinemuchi, K., Smith, H.A., Woźniak, P.R., McKay, T.A., ROTSE Collaboration, 2006, *AJ*, 132, 1202
- Kovács, G., Zucker, S., Mazeh, T., 2002, *A&A*, 391, 369
- Kozłowski, S., 2016, *ApJ*, 826, 118
- Lochner, M., McEwen, J.D., Peiris, H.V., Lahav, O., Winter, M.K., 2016, *ApJS*, 225, 31 (arXiv:1603.00882), doi:10.3847/0067-0049/225/2/31, ADS Link
- Lomb, N.R., 1976, *Astrophysics and Space Science*, 39, 447
- Long, J.P., El Karoui, N., Rice, J.A., Richards, J.W., Bloom, J.S., 2012, *PASP*, 124, 280
- Mahabal, A., Sheth, K., Gieseke, F., et al., 2017, arXiv e-prints, arXiv:1709.06257
- Malz, A.I., Hložek, R., Allam, J., T., et al., 2019, *AJ*, 158, 171
- Martínez-Palomera, J., Förster, F., Protopapas, P., et al., 2018, *AJ*, 156, 186
- McLaughlin, M.A., Mattox, J.R., Cordes, J.M., Thompson, D.J., 1996, *ApJ*, 473, 763, doi:10.1086/178188, ADS Link
- Mislis, D., Bachelet, E., Alsubai, K.A., Bramich, D.M., Parley, N., 2016, *MNRAS*, 455, 626 (arXiv:1511.03456), doi:10.1093/mnras/stv2333, ADS Link
- Mondrik, N., Long, J.P., Marshall, J.L., 2015, *ApJ*, 811, L34
- Moreno, J., Vogeley, M.S., Richards, G.T., Yu, W., 2019, *PASP*, 131, 063001
- Mowlavi, N., 2014, *A&A*, 568, A78
- Muthukrishna, D., Narayan, G., Mandel, K.S., Biswas, R., Hložek, R., 2019, *PASP*, 131, 118002
- Nandra, K., George, I.M., Mushotzky, R.F., Turner, T.J., Yaqoob, T., 1997, *ApJ*, 476, 70, doi:10.1086/303600, ADS Link

- Narayan, G., Zaidi, T., Soraisam, M.D., et al., 2018, *ApJS*, 236, 9 (arXiv:1801.07323), doi:10.3847/1538-4365/aab781, ADS Link
- Naul, B., Bloom, J.S., Pérez, F., van der Walt, S., 2018, *Nature Astronomy*, 2, 151
- Newling, J., Varughese, M., Bassett, B., et al., 2011, *MNRAS*, 414, 1987 (arXiv:1010.1005), doi:10.1111/j.1365-2966.2011.18514.x, ADS Link
- Nugent, P., Kim, A., Perlmutter, S., 2002, *PASP*, 114, 803 (arXiv:astro-ph/0205351), doi:10.1086/341707, ADS Link
- Nun, I., Protopapas, P., Sim, B., et al., 2015, arXiv e-prints, arXiv:1506.00010
- Pashchenko, I.N., Sokolovsky, K.V., Gavras, P., 2018, *MNRAS*, 475, 2326
- Protassov, R., van Dyk, D.A., Connors, A., Kashyap, V.L., Siemiginowska, A., 2002, *ApJ*, 571, 545
- Revsbech, E.A., Trotta, R., van Dyk, D.A., 2018, *MNRAS*, 473, 3969 (arXiv:1706.03811), doi:10.1093/mnras/stx2570, ADS Link
- Richards, J.W., Starr, D.L., Butler, N.R., et al., 2011, *ApJ*, 733, 10 (arXiv:1101.1959), doi:10.1088/0004-637X/733/1/10, ADS Link
- Richards, J.W., Homrighausen, D., Freeman, P.E., Schafer, C.M., Poznanski, D., 2012a, *MNRAS*, 419, 1121
- Richards, J.W., Starr, D.L., Miller, A.A., et al., 2012b, *ApJS*, 203, 32
- Saha, A., Vivas, A.K., 2017, *AJ*, 154, 231
- Sako, M., Bassett, B., Connolly, B., et al., 2011, *ApJ*, 738, 162 (arXiv:1107.5106), doi:10.1088/0004-637X/738/2/162, ADS Link
- Sánchez-Sáez, P., Reyes, I., Valenzuela, C., et al., 2021, *AJ*, 161, 141
- Scargle, J.D., 1982, *ApJ*, 263, 835
- Scargle, J.D., Norris, J.P., Jackson, B., Chiang, J., 2013, *ApJ*, 764, 167 (arXiv:1207.5578), doi:10.1088/0004-637X/764/2/167, ADS Link
- Schwarzenberg-Czerny, A., 1989, *MNRAS*, 241, 153
- Selam, S.O., 2004, *A&A*, 416, 1097
- Sokolovsky, K.V., Lebedev, A.A., 2018, *Astronomy and Computing*, 22, 28

- Sokolovsky, K.V., Gavras, P., Karamelas, A., et al., 2017, MNRAS, 464, 274
- Stellingwerf, R.F., 1978, ApJ, 224, 953
- Stetson, P.B., 1996, PASP, 108, 851, doi:10.1086/133808, ADS Link
- Stringer, K.M., Long, J.P., Macri, L.M., et al., 2019, AJ, 158, 16 (arXiv:1905.00428), doi:10.3847/1538-3881/ab1f46, ADS Link
- Süveges, M., 2012, arXiv e-prints, arXiv:1212.0645
- [LDM-151]**, Swinbank, J.D., et al., 2017, *Data Management Science Pipelines Design*, LDM-151, URL <https://ls.st/LDM-151>
- Tamuz, O., Mazeh, T., North, P., 2006, MNRAS, 367, 1521
- Torrealba, G., Catelan, M., Drake, A.J., et al., 2015, MNRAS, 446, 2251
- VanderPlas, J.T., 2018, ApJS, 236, 16
- VanderPlas, J.T., Ivezić, Ž., 2015, ApJ, 812, 18
- Varughese, M.M., von Sachs, R., Stephanou, M., Bassett, B.A., 2015, MNRAS, 453, 2848 (arXiv:1504.00015), doi:10.1093/mnras/stv1816, ADS Link
- Villar, V.A., Berger, E., Miller, G., et al., 2019, ApJ, 884, 83
- Villar, V.A., Hosseinzadeh, G., Berger, E., et al., 2020, ApJ, 905, 94
- Vincenzi, M., Sullivan, M., Firth, R.E., et al., 2019, MNRAS, 489, 5802 (arXiv:1908.05228), doi:10.1093/mnras/stz2448, ADS Link
- Wang, Y., Khardon, R., Protopapas, P., 2012, ApJ, 756, 67
- Welch, D.L., Stetson, P.B., 1993, AJ, 105, 1813, doi:10.1086/116556, ADS Link
- Wozniak, P.R., 2000, Acta Astronomica, 50, 421
- Wozniak, P.R., Udalski, A., Szymanski, M., et al., 2001, Acta Astronomica, 51, 175
- Zechmeister, M., Kürster, M., 2009, A&A, 496, 577
- Zucker, S., 2016, MNRAS, 457, L118

B Acronyms

Acronym	Description
ADQL	Astronomical Data Query Language
AGN	active galactic nuclei
ANTARES	Arizona-NOA Temporal Analysis and Response to Events System
AP	Alert Production
API	Application Programming Interface
B	Byte (8 bit)
DM	Data Management
DMTN	DM Technical Note
DPDD	Data Product Definition Document
DRP	Data Release Production
LDM	LSST Data Management (Document Handle)
LSE	LSST Systems Engineering (Document Handle)
LSST	Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope)
PPDB	Prompt Products DataBase
PSF	Point Spread Function
RFC	Request For Comment
SN	SuperNovae
TAI	International Atomic Time
TAP	Table Access Protocol
VLT	Very Large Telescope (ESO)
VO	Virtual Observatory
ZTF	Zwicky Transient Facility