Vera C. Rubin Observatory
Data Management

Periodicity Analysis in Alert Production

Anastasios Tzanidakis, and Eric C. Bellm

DMTN-221

https://doi.org/10.71929/rubin/2998097

Latest Revision: 2025-10-10

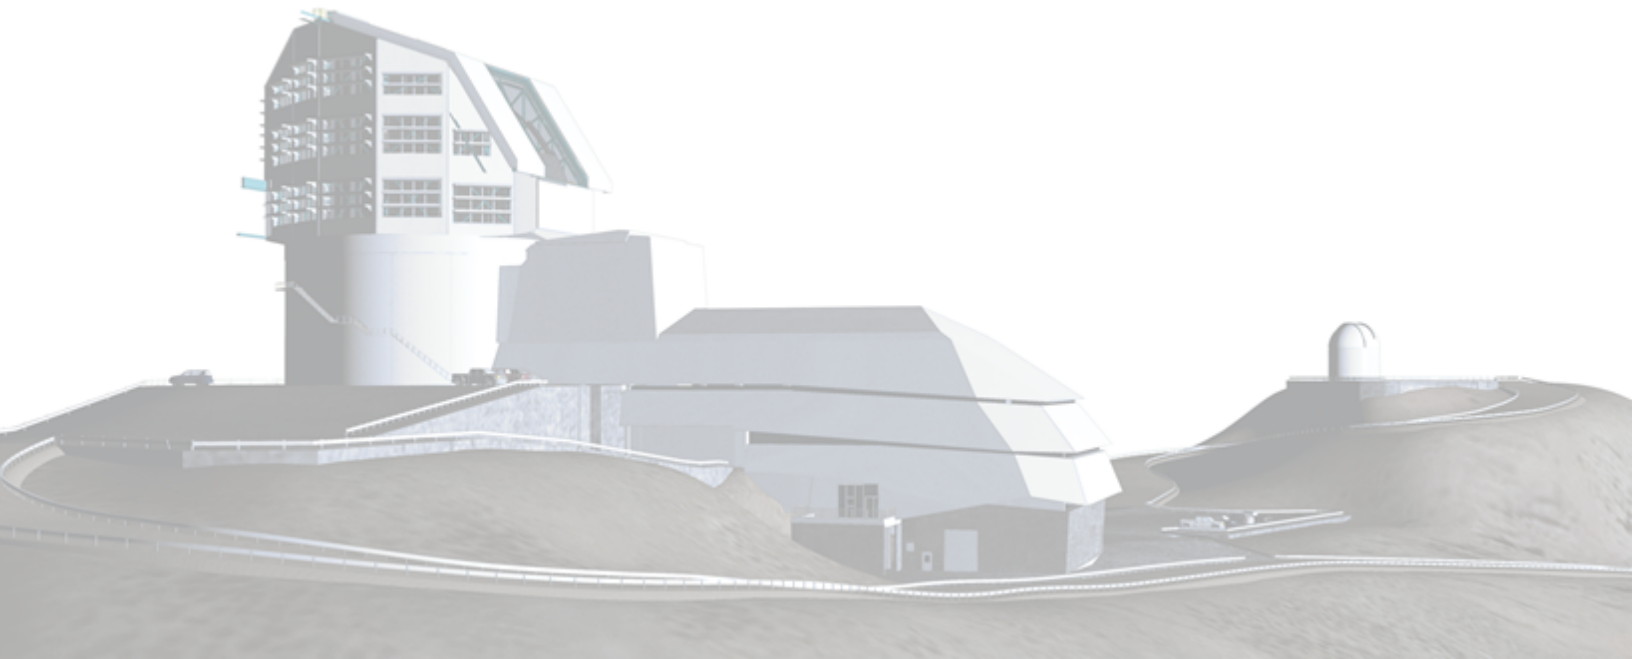

Abstract

The baselined timeseries features to be computed in Alert Production include a Lomb-Scargle periodogram ran on two classes of variable systems: RR Lyrae and Eclipsing Binaries. Based on a simulated LSST-like cadence light curves taken from the Extended LSST Astronomical Time-series Classification Challenge (ELAsTiCC) we perform an end-to-end test to characterize the periodicity recovery on the Alert Production multi-band light curves. In both variable classes, we found that a single-band Lomb-Scargle implementation yields to a low fraction of recovered periods, with a significant preference on the simple periodic phenomena such as RR Lyrae. We also investigated the results from a multi-band Lomb-Scargle and found an increased fraction of recovered periodicities above 15% for the eclipsing binaries, and over 80% for the RR Lyrae stars. Our findings suggest that a multi-band Lomb-Scargle should be implemented for searching periodic phenomena through AP. We also asses the computational and scientific performance of several configurations on simulated alert data and find that our current configuration scales linearly with the number of detections while assuming an heuristic frequency grid.

Change Record

Version	Date	Description	Owner name
1	2024-05-03	Unreleased.	Anastasios Tzanidakis, Eric Bellm
2	2025-10-10	Add DOI	Tim Jenness

Document source location: `https://github.com/lsst-dm/dmtn-221`

Contents

1 Motivation **1**

2 Synthetic Light Curves **1**

3 Injection-Recovery Testing **2**

3.1 Period Grid 3

3.2 Single Band Lomb-Scargle Periodogram 4

3.3 Multi-Band Lomb-Scargle Periodogram 5

3.4 Peak Significance Metric 6

3.5 Timing Analysis 9

3.6 Current Limitations 11

3.7 Implementation to LSST Stack 13

3.8 Conclusions 14

A References **14**

B Acronyms **15**

Periodicity Analysis in Alert Production

1 Motivation

The characterization of periodicity from time-series is a fundamental constraint to numerous astrophysical applications. Phenomenologically, estimating the periodicity and its significance can shed light on stellar pulsation theory (Antonello & Pastori, 1981), distance estimation and mapping of the Galaxy through the period-luminosity relationship (Skowron et al., 2019), constraint fundamental parameters of stellar binaries (Farinella et al., 1979) and stellar rotation (Walkowicz & Basri, 2013). In the recent decade, the use of periodicity has also been extensively used as a feature to classify variable phenomena (Richards et al., 2011) across the Hertzsprung-Russell diagram.

Previous work from Oluseyi et al. (2012) demonstrate the concept of injection-recovery test for synthetically generated RR Lyrae templates from SDSS using the op_sim version 1.29 cadence strategy. However, no study to date has demonstrated the recovery period distribution of more complex periodic phenomena such as eclipsing binaries in real-time. In this study, we focus specifically in the context of the LSST Alert Production system. The LSST Prompt pipelines that will be completed via the DIA will identify changing sources in the difference images. Alerts will be issued 60 seconds after the of the image readout. In this study, we perform an end-to-end injection-recovery analysis for periodic multi-band time series to asses the statistical significance of period finding in the context of the LSST AP as suggested in DMTN-118.

2 Synthetic Light Curves

In this study we use the training set Extended LSST Astronomical Time-series Classification Challenge (ELAsTiCC) light curves with a 12 month history that will best mimic the calculations performed for AP. In short, ELAsTiCC is a real-time pipeline for generating mock photometric alerts at the expected LSST rate. The photometric alerts will be distributed in real-time to brokers to benchmark classification algorithms. The ELAsTiCC dataset uses the v.1.7 cadence strategy from op_sims of the first three years of the survey. A single detection is based on the DIA performance from DC2. Each light curve include photometric noise using Poisson noise that includes the equivalent area, background noise per unit area, sky and CCD read

noise.Most importantly, the ELAsTiCC light curves simulate force photometry alerts which we recommend all time series features to be calculated upon.

For this investigation, we use two classes of periodic variable sources: RR Lyrae (RRL) and eclipsing binaries (EB). While other quasi-periodic sources, such as AGN, might be of importance, in this study we aim to answer if computing the periodogram on periodic phenomena is possible.We note one small drawback using the ELAsTiCC training is the relatively small sample size (N=360) and low signal-to-noise ratio for the eclipsing binaries. In Section 3.6 we discuss the implications of the SNR of the EB light curves.To maximize the use different eclipsing binary stars sampled at different phases, each light curve over a three year photometric history is sampled at some random time and ensures that the maximum baseline is less than 12 months. While a small EB sample size limits the search for a larger distribution in periodicities, it remains nonetheless useful to see the periodicity recovery rates at different sampling phases. In Figure 1, we show the underlying injected period distribution or both periodic classes.Finally, in Figure 2 we show the typical multi-band AP light curve sampled from ELAsTiCC with uncertainties.

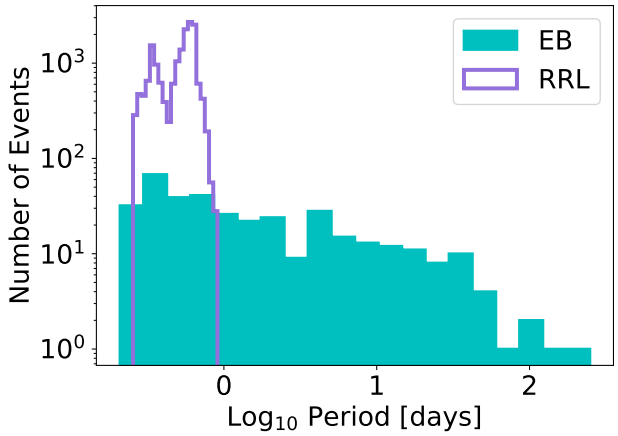

FIGURE 1: True period distribution from the ELAsTiCC training set.

3 Injection-Recovery Testing

In this section we discuss the periodicity recovery rate for injected RR Lyrae and eclipsing binaries in both scenarios of single-band and multi-band light curves. All computations of the periodogram henceforth will be the floating-mean Lomb-Scargle periodogram algorithm that is implemented from GATSPY VanderPlas & Ivezić (2015).

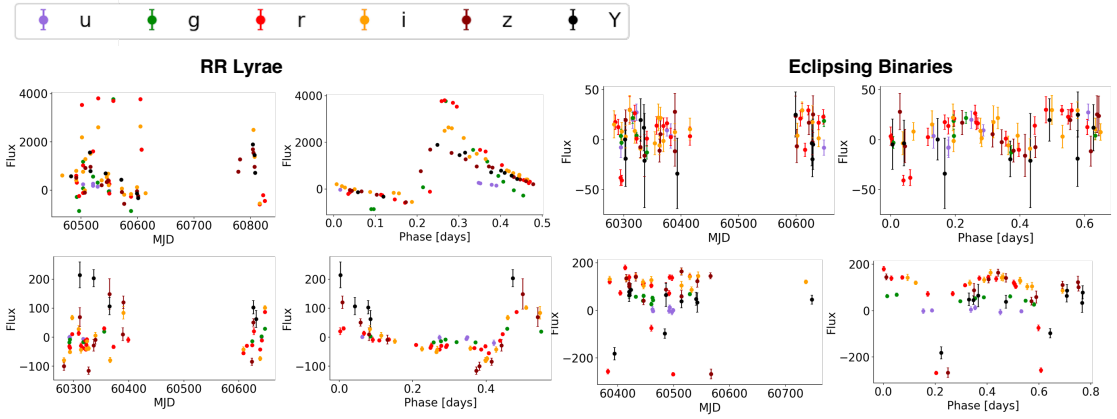

FIGURE 2: The above panels demonstrate the multi-band light curves simulated from ElasTiCC using RR lyrae (left panel) and eclipsing binaries (right panel). The first column of each panel shows the observed light curves, and the second the phase folded light curves at the correct period after 12 months.

3.1 Period Grid

The LSST will be search for periodic phenomena across both short and long timescales. The first challenge when computing periodograms in the context of AP is choosing the right trial period grid with separations that can lead to successful period finding.The periodogram is full of local minima and maxima, thus in practice the global maximum, which corresponds usually to the period with the highest power, is searched through a brute force period grid. High resolution period grids can also be very computationally expensive that is not feasible in the 60 second data latency of AP.

Inevitably, large searches in period space without an a priori will require an heuristic approach to generating an appropriate period grid that will account for the light curve characteristics. For this study, we use a heuristic frequency grid generator adapted in GATSPY that uses the light curve parameters to estimate the appropriate frequency grid length and separation. Unfortunately, for sparsely sampled light curves there does not exist a straightforward approach to finding the optimum frequency grid. It is inevitable that some light curves will either require coarser or finer frequency grids.

Generally, the frequency separation that we will be using for thus study can be approximated using the duration of the light curve, including some sampling constant ($\mathcal{O}$) and the total number of computed periods as a function of the Nyquist factor ($F_{nyquist}$), and number of detections (N_{det}):

$$\delta f = \frac{1}{(t_{max} - t_{min}) \times \mathcal{O}} \tag{1}$$

$$N_P = int\left(\frac{1}{2}\mathcal{O}F_{nyquist}N_{det}\right) \tag{2}$$

Building some intuition, the Nyquist factor pushes the minimum starting frequency to smaller periods while the oversampling factor will increase the maximum frequency grid and overall increase the resolution of the grid. In later sections we evaluate for each frequency grid the period accuracy computed on the ELAsTiCC light curves.

3.2 Single Band Lomb-Scargle Periodogram

We first consider the case of the canonical single-band detections from AP. In this section we aim to answer if the 12 month AP history in a single-band light curve is enough to recover periods.The reader should refer to VanderPlas & Ivezić (2015) for a throughout review of the Lomb-Scargle periodogram.

We consider the canonical floating-mean Lomb-Scargle to estimate the periodogram of 10^4 alert light curves in the *r*-band for three Fourier component modes. We caution the average LSST field, the completeness for single-band periodic events will likely depend on the filter since overall filter coverage is not uniform. For example, we would expect that the *u*-band filter to suffer from lower completeness due to the smaller number of visits. In Figure 3 we show the results of this analysis for RRL's (rop row) and EB's (bottom row). For each light curve model, we randomly sample a light curve and compute the Lomb-Scargle periodogram using a heuristic period grid with an oversampling factor of 2 and Nyquist frequency of 30 (~0.03 days) to account for the fast periodic phenomena. We extrapolated the highest power period from the periodogram and compare it to the true period. In Figure 3 we also include the number of *r*-band detections for each light curve. We note that we did not find any reliable periods beyond 5 days for the eclipsing binaries. This is likely due to the fact of the small sample size beyond 5 days. Since we are re-sampling the same underlying models it is possible that the light curve re-sampling is not adequate enough to identify correct large periods. As expected, we did not find any preference of an optical single-band photometric filter that outperformed the others. While the RR Lyrae exhibit at larger Fourier modeling modes more aliasing, overall the recovery fraction of correctly identified periods is larger compared to the EBs. For both

light curve models, we find that the single-band Lomb Scargle recovery fraction average to be higher for light curves that had more detections. Overall we find consistent results for the RR Lyrae done by a similar analysis by VanderPlas & Ivezić (2015). The true challenge of this approach comes from the more complex light curve models such as the EBs.

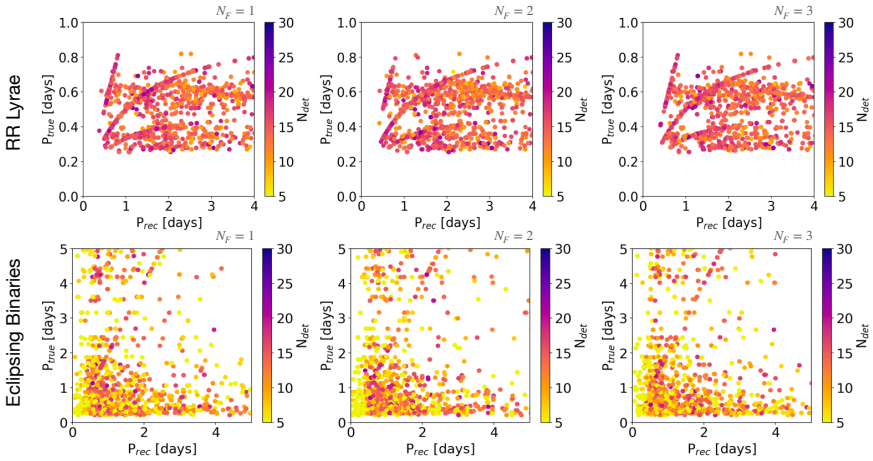

FIGURE 3: We show the injected (true period) versus recovered highest peak in the periodogram for a sample of 1000 *r*-band detections from RR lyrae and eclipsing binaries. Each column shows the period recovery for N=1 to N=3 Fourier components when computing the Lomb-Scargle periodogram.

On average, there will be approximately 80 detections in all photometric bands per 12 month light curve history. Given the sampling rate of LSST per pointing, this entails that the majority of single-band AP light curves would be too undersampled in order to compute a high confidence periodicity. Even with low completeness, computing the Lomb-Scargle periodogram on single-band time series alerts, LSST will likely find more periodic sources compared to any other known survey due to the large survey volume.The single-band Lomb-implementation are also faster by a factor of 5 due to the lower average number of detections per filter. We conclude that the single-band Lomb-Scargle misses by a higher margin the underlying true period due the small light curve photometric history. Now we turn our attention to the multi-band Lomb-Scargle.

3.3 Multi-Band Lomb-Scargle Periodogram

Next, we test the multi-band Lomb-Scargle periodogram implementation adapted in GATSPY to our ELAsTiCC light curves. Given the rich multi-band alert detections LSST will deliver, we expect that that conventional period finding algorithms to do better due to an increase of

number of total detections. In short, the multi-band Lomb-Scargle periodogram we will be considering for this study can be written out with two components that use implemented in GATSPY:

$$y = \theta_0 + \sum_{n=1}^{M_{base}} \left[\theta_{2n-1} \sin(n\omega t) + \theta_{2n} \cos(n\omega t)\right] \theta_0^{(k)} + \sum_{n=1}^{M_{band}} \left[\theta_{2n-1}^{(k)} \sin(n\omega t) + \theta_{2n}^{(k)} \cos(n\omega t)\right] \quad (3)$$

First is the base Fourier component that describes the overall shared variability, and second is the per band component that is modeled by the residuals of the base component with each photometric filter. In practice, a larger configuration of base and band components can lead to the estimation of more complex light curves. GATSPY automatically determines through the use of least-squares minimization the scalar parameters from the above equation. For a more detailed mathematical review of the multi-band Lomb-Scargle we suggest the review of VanderPlas & Ivezić (2015). In 4 we show the overall increase of performance in recovery fraction from the multi-band Lomb-Scargle. Particularly, we notice that the RRL are over 50% complete across all bins of number of total detections, and almost 99% complete for the highest number of detections bin. On the contrary, we see that the EB mutli-band completeness does not outcompete the single-band approach especially for the smaller detection bins. We only see a larger completeness performance for EB light curves with more than 70 detections across all filters.

3.4 Peak Significance Metric

One imminent challenge beside choosing the right period grid, and model complexity is the reported period significance. It is inevitable that the AP periodic pipeline will likely report spurious periods, and both harmonics and aliases of the underlying period of the source. The effects of aliasing and harmonics are typically identified by visual inspection of the phase folded light curves. In this section we explore empirical metrics to evaluate peak significance within the AP time frame. One common diagnostic used in the literature is the False Alarm Probability (FAP) a given peak in the power spectrum. The FAP measures the probability uncertainty associated with the selected peak under the null hypothesis. Generally, the reported FAP for a specific period peak can be used to inform the user of its a trustworthy detection.

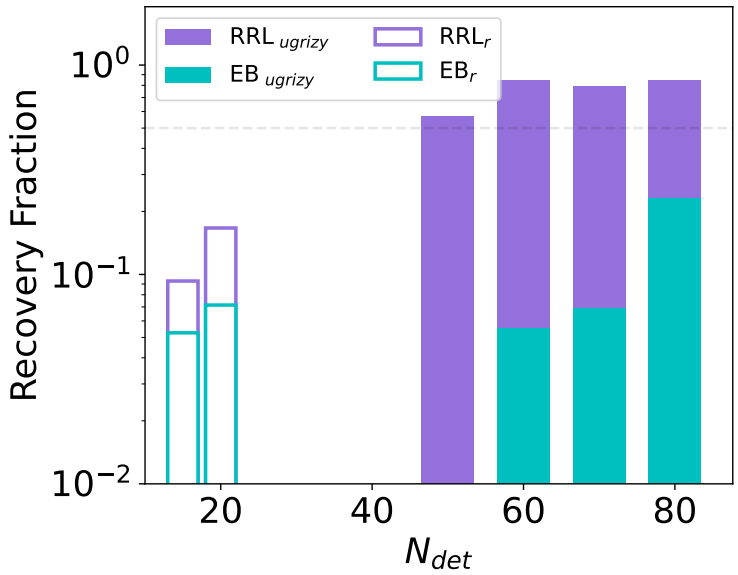

FIGURE 4: Bar graph of the recovery fraction within 0.001 days of the period as a function of number of detections. The lined bars are the results from the single-band Lomb-Scargle periodogram assuming a two Fourier components. The filled bar graphs are the multi-band Lomb-Scargle periodogram assuming a (N_{base}, N_{band})=(2,1) model configuration.

Given the absence of a true analytic solution to the Lomb-Scargle FAP, most methods rely on computational approaches such as bootstrapping. Bootstrapping by virtue is computationally expensive. For example, a 1% FAP rate would require 1000 bootstraps per light curve. By simply scaling the expected number of periodic sources to be found in each LSST pointing, we find that that the number of computed periodogram to likely exceed the computing resources of AP. Given the fast AP latency and number of periodic sources, a significant real-time bootstraping is likely not feasible for AP. A more popular approach is to use parametric FAP estimators that asymptotically converge to similar distributions of bootstraping FAP approaches. A sophisticated approach suggested by Baluev (2008) which uses extreme value statistics to compute an upper bound of the false alarm probability for the alias-free case:

$$P_{FAP} \approx \frac{\Gamma((N-1)/2)}{\Gamma((N-2)/2)} \sqrt{4\pi \,\mathrm{var}\,(t_i)}\, f_{\mathrm{u}}\, (1 - z_{\mathrm{obs}})^{\frac{N-4}{2}} \sqrt{z_{\mathrm{obs}}} \tag{4}$$

Where z_{obs} is the power of the recovered peak and f_u the maximum computed frequency for the periodogram. The Baluev approximation has already been applied in other studies such as Süveges et al. (2015) and is suggested as a computationally inexpensive and reliable form to approximate the FAP, especially in the case of low aliasing and spectral leakage. More ro-

bust techniques that account for correlated noise should be considered in future studies that want to extrapolate FAP for correlated noise measurements such as the technique presented by Delisle et al. (2020). For this study we will only consider the Baluev approximation. The more complex issue arrises from the complex mathematics that involve the derivation of the Baluev approximation for the multi-band Lomb-Scargle form. The derivation of such form is beyond the scope of this paper, and we suggest future studies to explore if such solution exists. For practical purposes, in our case we assume that the Baluev (2008) form will hold true for the multi-band case, however we caution that the overall FAP estimates to possibly be underestimated or overestimated.

To test whether the Baulev approximation holds true for the multi-band light curves we attempt to correlate a conventional bootstraping significance approach with the Baluev form. We begin with a 500 light curve sample of RRL and EB assuming a (N_{base}, N_{band})=(2,1) model configuration. For each light curve, we run the multi-band Lomb-Scargle implementation 1000 times by sampling the light curve without replacement and extrapolate the maximum power for each iteration. Finally, we compute the ratio between the recovered power of the original Lomb-Scargle periodogram and the median bootstrapped power, which can be interpreted as a significance estimator. In Figure 5, we show the correlation between the Baluev and bootstrap approach for each variable class and color-coded by the number of *ugrizy* detections. For each variable class, we ran a Pearson-R correlation test and found in both cases a a negative r-slope with p-value significance p$\ll$ 0.005. We also notice that the overall FAP levels of the EBs to be higher since their recovery rates are much smaller.

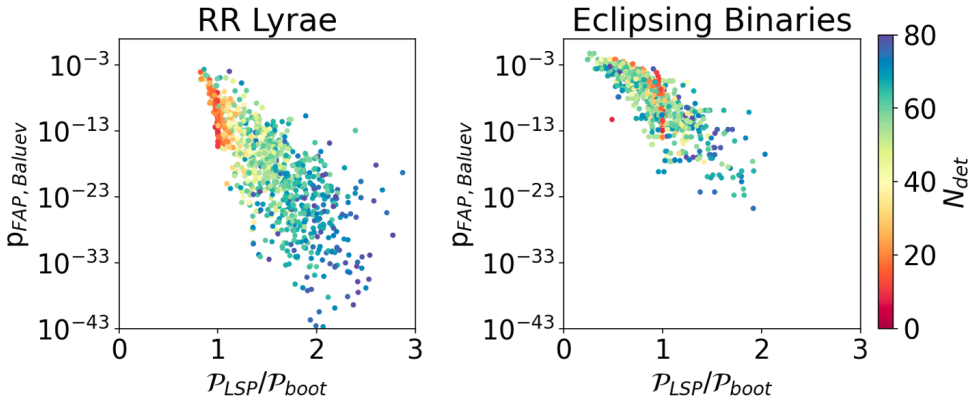

FIGURE 5: Multi-band FAP rate computed via the Baluev approximation compared to a 2σ bootstrap approach for 5000 RRL and EB alert light curves. For each light curve we also color code the total number of $ugrizy$ detections.

Based on preliminary estimates, Figure 5 suggest that parametric forms of the FAP can approx-

imate the bootstrap significance without the computational cost of running N bootstraps. In the cases of high SNR such as the RR Lyrae, we find that both significance estimates show a dependency of the number of detections.

3.5 Timing Analysis

In this section we bench mark the average run time of the Lomb-Scargle periodogram implementation from GATSPY. The balancing act of high periodocity accuracy while confining the calculations within the AP timeframe is a challenging act. While coarse period grids and complex models can achieve more accurate periods, this will not be possible to do with AP. Instead, in this section we will explore the lower limit of run time while tweaking the overall period grid coarseness and model complexity.

The current GATSPY, which is a pure Python based implementation of the Lomb-Scargle periodogram, scales as $\mathcal{O}[N]$ algorithm for small N and becomes $\mathcal{O}[N^2]$ at larger trial periods.The unique rapid calculation of the periodogram emerges from the efficient use of linear algebraic operations. The average run time also depends on the number of detections and model complexity which we will explore in this section. For our run time analysis we will consider both a single-band and multi-band case for the RR Lyrae only since the run time is not a function of the underlying signal.

We first measure the effects of run time on the single-band RRL photometry in the *r*-band filter. For each light curve, we run the Lomb-Scargle periodogram 10 times per light curve, and compute the median run time. Additionally, we run this for each combination of oversampling factor, Nyquist frequency, and Fourier components. In Figure 6, we display the results for the single-band photometry. We also color code the average fractional error per light curve iteration. As mentioned in the single-band Lomb-Scargle method, by design will be poorly sampled especially within the 12 month history. As expected, due to the sparse sampled light curves, the majority of light curves scale linearly with number of detections, and have average run times below one second. We do not find any major derivations from the average run time when increasing the number of Fourier components from one to three. It is possible that single-band periodograms will generally be able to afford higher modes of Fourier components while still fitting within the AP time frame. Overall the majority of single-band periodograms are consistently below one second for the suggested period grid parameters.

In Figure 7 we show the average run time for the multi-band light curves by increasing Fourier

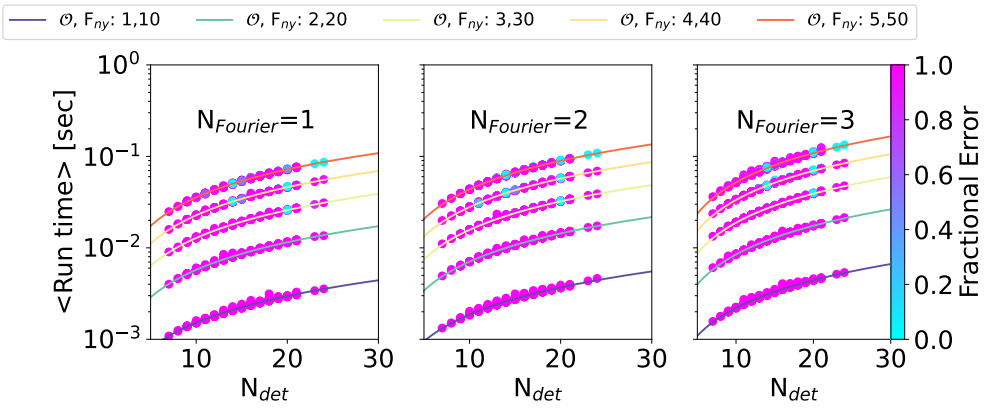

FIGURE 6: Average run time per number of r light curve detections. Each column represents the model complexity of each multi-band Lomb Scargle periodogram. Within each panel, we use a heuristically determined period grid by varying the oversampling factor and the the Nyquist frequency. Finally, for each successive run we color code the fractional error compared from the true period.

modes. Similarly to the single-band, we find a linear scaling between each period grid mode. We find that the finer period grids are better at identifying the correct periods while also keeping the run time at about one second per light curve with a 12 photometric history. It is evident that the multi-band Lomb-Scargle implementation has the advantage of similar run-times all while utilizing all six photometric bandpasses to find the period.

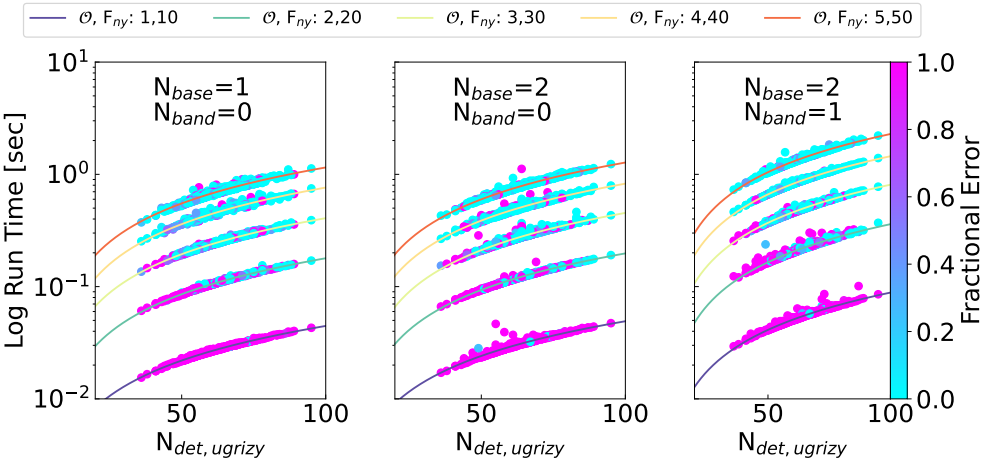

FIGURE 7: Average run time per number of $ugrizy$ light curve detections. Each column represents the model complexity of each multi-band Lomb Scargle periodogram. Within each panel, we use a heuristically determined period grid by varying the oversampling factor and the the Nyquist frequency. Finally, for each successive run we color code the fractional error compared from the true period.

Further tests should be done to explore the effects of the number density of periodic sources

per LSST pointing including other parameters such as achieved depth. Other metrics, such as minimum period grid required per pointing to recover a period source should also be investigated.

3.6 Current Limitations

From our analysis, it is evident that the eclipsing binaries overall lack completeness in both cases of single and multi-band. We suspect that the main culpruit of low completeness to be the average low SNR that computed for the ELAsTiCC eclipsing binary sample. For example, in Figure 8, we show for a sample of all training ELAsTiCC light curves with their average SNR for the RRL versus EB's. On average we found that the EBs are skewed toward a low SNR while many more RRL have high amplitude variations and thus higher SNR. Given the noisy EB light curves provided from ELAsTiCC is is possible that our current estimates would be considered for periodic events close to the limiting LSST magnitude.

As a proof of concept, we can validate a possible scenario of high SNR eclipsing binary light curves using the cadence strategy of op_sims v.1.7 using an in-house light curve building tool. First, we will assume that the phase-folded EB light curve can be modeled by a decomposed four-component Fourier series model:

$$m(t) = m_0 + \sum_{i=1}^{N=4} A_i cos(2\pi i\Phi) + B_i sin(2\pi i\Phi) \tag{5}$$

where the phase-offset can be defined as:

$$\Phi(t, t_0) = \frac{t - t_0}{P} - int\left(\frac{t - t_0}{P}\right) \tag{6}$$

We adopt the fitted Fourier terms from Deb & Singh (2011) where the parameters for 62 luminous bright single-band eclipsing binary light curves are modeled from the All Sky Automated Survey (ASAS)-3 project. We assume that the Fourier terms apply to all photometric bandpasses, while adding a small color offset between each photometric filter (by roughly assuming that these are equal temperature eclipsing binaries). Using the 62 EB models, we sample each model at the $ugrizy$ cadence separation described by the op_sims v.1.7 simulation. For each EB event, we randomly assign a coordinate within the LSST footprint and a distance drawn from a log-normal distribution with mean 5 kpc and standard deviation of 2.

Using a log-normal distributon of absolute magnitudes we then solve the distance modulus to infer the apparent magnitude of the system, without taking into account the effects of Galactic extinction. Additionally, for each photometric bandpass we assume a random constant color offset between zero and 4 magnitudes. To estimate the expected SNR from a source with magnitude m by assuming a Gaussian distribution of photons we apply the standard $m5snr$ equation. Finally, to compute the photometric error we use the m_5 SNR estimate the photometric errors associated with each light curve. We run the multi-band Lomb-Scargle on the multi-band light curves and asses the recovery fraction. In Figure 9 we show the injected versus recovered period for our custom light curves with overall higher SNR compared to the ELAsTiCC light curves overall recover a larger fraction the true underlying injected periods. We notice that the completeness especially increases at the (N_{base}, N_{band})=(2,1) Fourier model configuration.

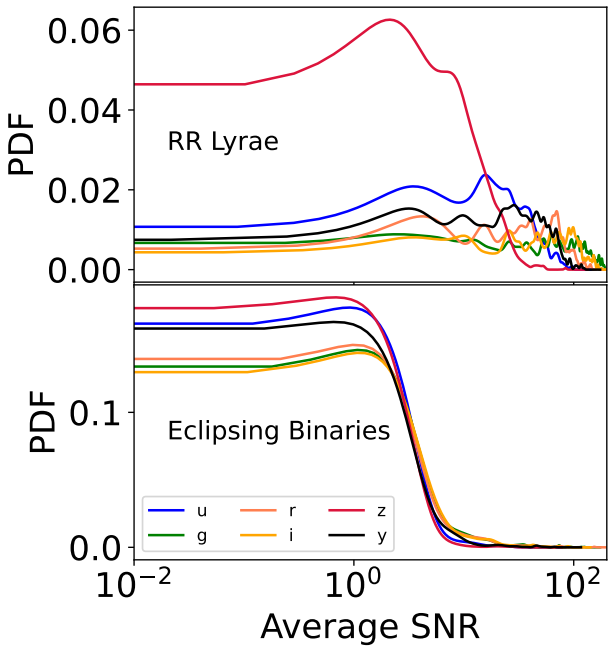

FIGURE 8: Distribution of signal-to-noise ratio for simulated ELAsTiCC light curves. We compute the median signal-to-noise ratio per light curve in the $ugrizy$ bandpasses. We smooth each distribution using a Gaussian Kernel Density estimator. Top panel is the RR Lyrae and bottom is the Eclipsing Binaries.

Future investigations prior to the onset of the LSST should investigate a larger sample of eclipsing binary populations at higher luminosities and more throughout combinations of physical parameters.

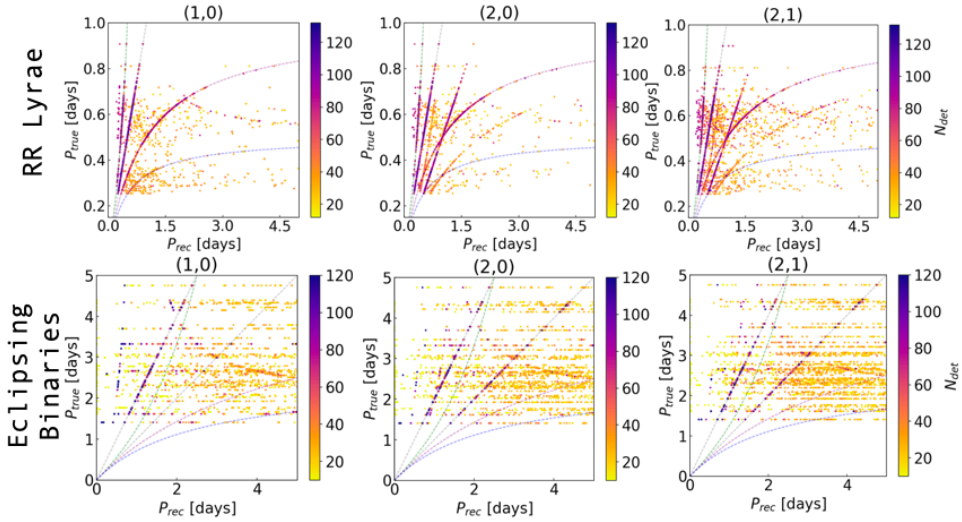

FIGURE 9: Multi-band injection-recovery test for 5000 RRL (top row) and EB (bottom row). Each column title includes the number of Fourier base and band terms used to compute the Lomb-Scargle periodogram. Additionally we color code each recovered period by the number of detections. The overlaid curved lines represent the n=1 aliasing, while the straight lines represent the n=1,2 harmonics.

3.7 Implementation to LSST Stack

As of 2023, `Astropy`[1] has implemented for public use the multi-band Lomb-Scargle periodogram directly from the GATSPY implementation. To use the ASTROPY MLSP implementation, we followed the online open-source documentation tutorial for the basic use of the API that has a slightly different API from the inherent GATSPY implementation[2]. Using the same ELAsTiCC light curves and period grid parameters as in the previous sections, we found similar performance both in timing and period recovery between both algorithms. The current implementation on the Stack assumes (N_{base}, N_{band})=(1,1), and significance is evaluated via the Baluev approximation (i.e see equation 3 in Süveges et al. (2015)). As part of releasing the multi-band Lomb-Scargle period and its significance, we also report the model fit parameters described in equation 3. Assuming each light curve as (k) unique bands, we define the Amplitude (A^k) and Phase (ϕ^k) from the terms in the linear model for each filter:

$$A^k = \sqrt{(\theta_1^k)^2 + (\theta_2^k)^2} \tag{7}$$

$$\phi^k = arctan(\theta_2^k/\theta_1^k) \tag{8}$$

[1] ASTROPY v5.3.3
[2] https://docs.astropy.org/en/stable/timeseries/lombscarglemb.html

3.8 Conclusions

In conclusion, this study investigates the periodicity recovery rate of RR Lyrae and eclipsing binaries on single-band and multi-band light curves using the Extended LSST Astronomical Time-series Classification Challenge (ELAsTiCC) training dataset. Our results suggest that using a fixed heuristic frequency grid can successfully recover the true underlying periods from the Lomb-Scargle periodogram using GATSPY. We ran single-band Lomb-Scargle on the *r*-band light curves and overall found our recovery fraction to be less than 30% for RRL and less than 10% for EB. On the contrary, we find a significant leverage and boost of performance using the multi-band Lomb-Scargle approach and find the RRL recovery fraction to be as high as 95 %, while the EBs to be 40 %. We speculate that the low recovery fraction of the EB's are likely due to the injected low SNR models in the ELAsTiCC training dataset. After using an in-house light curve generation tool using the decomposed Fourier terms for EB's and RRL, we find an improvement in the overall recovery fraction. This work also suggest that the analytic Baluev FAP approximation correlates significantly with a more traditional bootstrap FAP approach, hence we suggest that this can be used as a proxy for peak significance. Within the time limitations in AP, we demonstrate through a benchmark timing analysis that a heuristic frequency for both single-band and multi-band Lomb-Scargle periodograms with a 12-month photometric history are scalable to the time requirements of the LSST AP. It is suggested that future investigations prior to the onset of the LSST AP should investigate a larger suite of eclipsing binary populations at higher luminosities and with more throughout combinations of physical parameters, including a more robust analysis on the FAP estimates for cases of correlated noise and high aliasing.

A References

References

Antonello, E., Pastori, L., 1981, PASP, 93, 237, doi:10.1086/130812, ADS Link

Baluev, R.V., 2008, MNRAS, 385, 1279 (arXiv:0711.0330), doi:10.1111/j.1365-2966.2008.12689.x, ADS Link

[DMTN-118], Bellm, E.C., 2023, *Review of Timeseries Features*, Data Management Technical Note DMTN-118, NSF-DOE Vera C. Rubin Observatory, URL https://dmtn-118.lsst.io/, doi:10.71929/rubin/2997999

Deb, S., Singh, H.P., 2011, MNRAS, 412, 1787 (arXiv:1011.2574), doi:10.1111/j.1365-2966.2010.18016.x, ADS Link

Delisle, J.B., Hara, N., Ségransan, D., 2020, A&A, 635, A83 (arXiv:2001.10319), doi:10.1051/0004-6361/201936905, ADS Link

Farinella, P., Luzny, F., Mantegazza, L., Paolicchi, P., 1979, ApJ, 234, 973, doi:10.1086/157581, ADS Link

Oluseyi, H.M., Becker, A.C., Culliton, C., et al., 2012, AJ, 144, 9, doi:10.1088/0004-6256/144/1/9, 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

Skowron, D.M., Skowron, J., Mróz, P., et al., 2019, Science, 365, 478 (arXiv:1806.10653), doi:10.1126/science.aau3181, ADS Link

Süveges, M., Guy, L.P., Eyer, L., et al., 2015, MNRAS, 450, 2052 (arXiv:1504.00782), doi:10.1093/mnras/stv719, ADS Link

VanderPlas, J.T., Ivezić, Ž., 2015, ApJ, 812, 18 (arXiv:1502.01344), doi:10.1088/0004-637X/812/1/18, ADS Link

Walkowicz, L.M., Basri, G.S., 2013, MNRAS, 436, 1883 (arXiv:1309.2159), doi:10.1093/mnras/stt1700, ADS Link

B Acronyms

Acronym	Description
AGN	Active Galactic Nuclei
AP	Alert Production
API	Application Programming Interface
B	Byte (8 bit)
CCD	Charge-Coupled Device
DC2	Data Challenge 2 (DESC)
DIA	Difference Image Analysis
DMTN	DM Technical Note

DOI	Digital Object Identifier
ELAsTiCC	Extended LSST Astronomical Time Series Classification Challenge
LSST	Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope)
RRL	RR Lyrae stars
SDSS	Sloan Digital Sky Survey
SNR	Signal to Noise Ratio