The Astronomical Journal, 168:181 (21pp), 2024 October https://doi.org/10.3847/1538-3881/ad6cdd

Asteroid Period Solutions from Combined Dense and Sparse Photometry

Michael Gowanlock[1,2], David E. Trilling[1,2], Andrew McNeill[3,4], Daniel Kramer[1,2], and Maria Chernyavskaya[2]
[1] School of Informatics, Computing, and Cyber Systems, Northern Arizona University, P.O. Box 5693, Flagstaff, AZ 86011, USA; michael.gowanlock@nau.edu
[2] Department of Astronomy and Planetary Science, Northern Arizona University, P.O. Box 6010, Flagstaff, AZ 86011, USA
[3] Department of Physics, Lehigh University, 16 Memorial Drive East, Bethlehem, PA 18015, USA
[4] Department of Physics and Astronomy, Bowling Green State University, Bowling Green, OH 43403, USA
Received 2023 December 22; revised 2024 July 13; accepted 2024 August 5; published 2024 September 30

Abstract

Deriving high-quality light curves for asteroids and other periodic sources from survey data is challenging owing to many factors, including the sparsely sampled observational record and diurnal aliasing, which is a signature imparted into the periodic signal of a source that is a function of the observing schedule of ground-based telescopes. In this paper we examine the utility of combining asteroid observational records from the Zwicky Transient Facility and the Transiting Exoplanet Survey Satellite, which are the ground- and space-based facilities, respectively, to determine to what degree the data from the space-based facility can suppress diurnal aliases. Furthermore, we examine several optimizations that are used to derive the rotation periods of asteroids, which we then compare to the reported rotation periods in the literature. Through this analysis we find that we can reliably derive the rotation periods for ∼85% of our sample of 222 objects that are also reported in the literature and that the remaining ∼15% are difficult to reliably derive, as many are asteroids that are insufficiently elongated, which produces a light curve with an insufficient amplitude and, consequently, an incorrect rotation period. We also investigate a binary classification method that biases against reporting incorrect rotation periods. We conclude the paper by assessing the utility of using other ground- or space-based facilities as companion telescopes to the forthcoming Rubin Observatory.

Unified Astronomy Thesaurus concepts: Asteroids (72); Astroinformatics (78); Light curves (918); Small Solar System bodies (1469); Sky surveys (1464)

Materials only available in the online version of record: machine-readable table

1. Introduction

There are several benefits of using either ground- or space-based telescopic surveys to derive the physical properties of astrophysical phenomena, and in this paper we examine the physical properties of asteroids.

It is well-known that when deriving the period for periodic sources ground-based observatories will impart a periodic signal as a function of the diurnal cycle (J. T. VanderPlas 2018). These aliases typically appear in a periodogram at frequencies that are multiples of 1 day^{-1}. In contrast, space-based observatories do not suffer from diurnal aliases because they are not prevented from observing at any time during the day.

The Transiting Exoplanet Survey Satellite (TESS) is a space-based observatory with a primary mission to detect transiting exoplanets (G. R. Ricker et al. 2015), where the telescope stares at a field for a long duration. This telescope has observed intervening asteroids within its field of view. TESS rarely encounters the same asteroid between two pointings, which limits the temporal extent for which a given object is observed. This time window directly impacts the range of viable rotation periods for an asteroid. In contrast, ground-based surveys are able to derive the rotation periods of asteroids with long periods because they may observe the asteroid numerous times over several years.

In summary, in the context of surveys, ground-based telescopes are limited by diurnal observing schedules, and space-based telescopes that stare at a field for long durations are only sensitive to detecting asteroids with short rotation periods. By combining data sets from ground- and space-based observatories, there are several potential benefits, including (i) decreasing the fraction of objects that have a period solution at aliases, (ii) improving the sensitivity to long rotation periods, and (iii) utilizing partial light-curve data that are of little value on their own but could be used to augment the observational records of other catalogs.

Based on the above motivation, in this paper we examine two applications for combining data from ground- and space-based observatories, which are summarized as follows:

1. To improve the overall fidelity of derived asteroid rotation period solutions by comparing to the rotation periods reported in the literature.
2. To examine the use of a companion telescope (either ground or space based) to the Rubin Observatory to assess how this companion may improve rotation period fidelity.

In the context of surveys, periodic sources often have light-curve periods that are difficult to constrain because surveys produce sparse photometry. The first application above is similar to other papers in the literature that develop new methods that either maximize the fraction of correctly assigned light-curve periods in survey data or examine which methods should be applied to a particular application (A. J. Drake et al. 2013; M. J. Graham et al. 2013; M. Süveges et al. 2015; R. J. Oelkers et al. 2017; M. W. Coughlin et al. 2021; D. Kramer et al. 2023b). The second application is motivated by assessing the impact that a companion telescope to the Rubin Observatory would have on improving solar system science.

The paper is organized as follows. Section 2 outlines the data sets used in the paper. Section 3 describes the period-finding algorithm used in the paper and associated parameters. It also illustrates the baseline distribution of rotation periods across each

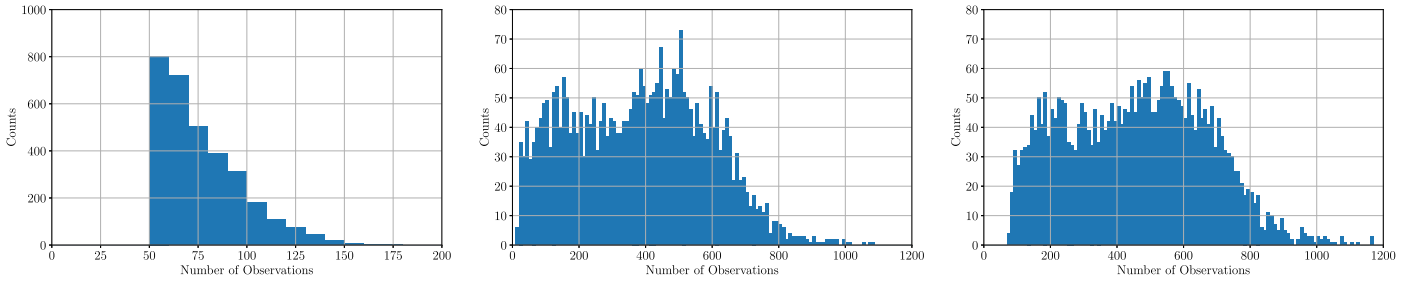

(a) SS1 ($\mu = 76.1$) (b) TESSC1 ($\mu = 382.8$) (c) SS1∩TESSC1 ($\mu = 458.9$)

Figure 1. The distribution of the number of observations for the objects in the (a) SS1, (b) TESSC1, and (c) SS1∩TESSC1 data sets, limited to the $n = 3168$ objects appearing in both SS1 and TESSC1. The bin width is 10 observations in each histogram.

Table 1
Summary of Data Sets and Properties

Data Set Name	Description	Number of Objects	Number of Observations
SS1	The first SNAPS alert broker data release containing light-curve properties of asteroids derived from ZTF (D. E. Trilling et al. 2023b).	28,638	2,145,478
TESSC1	Light-curve properties of asteroids derived from TESS (A. McNeill et al. 2023).	28,878	5,637,892
SS1∩TESSC1	The intersection of objects in SS1 and TESSC1.	3168	1,453,748

data set. Section 4 presents the proposed optimizations used to maximize derived period fidelity as compared to the solutions in the literature that are reported by the Light Curve Data Base (LCDB). With a selection of optimizations established, Section 5 applies them to the full sample of objects that are found in both Zwicky Transient Facility (ZTF) and TESS data sets. Section 6 assesses the utility of using other facilities to augment the capabilities of the Rubin Observatory. Finally, Section 7 concludes the paper.

2. Data

Table 1 outlines the data sets used in this paper. SS1 and TESSC1 are from ZTF (E. C. Bellm et al. 2019) and TESS (G. R. Ricker et al. 2015), respectively. SS1∩TESSC1 contains only those objects that are found in both SS1 and TESSC1.

As will be shown, we compare our period solutions to LCDB (B. D. Warner et al. 2009). However, there are only 222 objects that are in SS1∩TESSC1 and LCDB (hereafter denoted as LCDB∩SS1∩TESSC1). While there are more objects in SS1∩TESSC1 (3168), we have no baseline for comparison for all of these objects, so we use the LCDB∩SS1∩TESSC1 data set for several of our comparisons.

We use a version of LCDB from 2020 that excludes TESS data deposited by A. Pál et al. (2020) and A. McNeill et al. (2023) such that when comparing our derived period solutions to LCDB we do not accidentally compare to some of our own period solutions as derived by our prior work in A. McNeill et al. (2023). Otherwise, we may inadvertently increase the total fraction of correctly derived period solutions.

2.1. Data Sets: SS1 and TESSC1

The SNAPShot1 (SS1) data set (D. E. Trilling et al. 2023b) contains ZTF (E. C. Bellm et al. 2019) observations of small bodies, with the vast majority being main-belt asteroids. In the SS1 data release, only those objects with ⩾51 observations have an assigned derived rotation period. This observation threshold was selected because a sufficient number of observations are needed to reliably derive a rotation period. Of those objects with ⩾51 observations, there are a total of 2,145,478 observations across 28,638 objects. A real/bogus score (D. A. Duev et al. 2019) was used to discard observations with high uncertainties (e.g., nonpoint sources and observations with poor subtractions were removed using this method). For more information on data processing see D. E. Trilling et al. (2023b).

TESSC1 contains 28,878 objects with a total of 5,637,892 observations, where the data were σ clipped to remove outliers in the observational record for each object (for more details on data processing see A. McNeill et al. 2023). Unlike the ZTF data in SS1, light curves were derived for all objects regardless of the number of observations (A. McNeill et al. 2023). Instead of using a threshold of 51 observations for SS1, A. McNeill et al. (2023) assigned a confidence score to each object that described the probability that the derived rotation period is correct.

Figures 1(a) and (b) report the distribution of the number of observations for each object in the SS1 and TESSC1 data sets, respectively, limited to those objects that are common to both data sets. Because the goal of this paper is to combine the two data sets together, we do not enforce a lower limit number of observations required of each object in the TESSC1 data set, as even a few observations in TESSC1 for an object may help supplement the observational record for SS1. However, as shown in Figure 1(b), the mean number of observations per object in TESSC1 is $\mu = 382.8$, and so there are a small fraction of objects that have a few observations.

2.2. Combining SS1 and TESSC1

We combine the two data records for SS1 and TESSC1 by taking the intersection of the objects in both data sets. This yields a total of $n = 3168$ objects, and as described above, we denote this data set as SS1∩TESSC1. We calculate H magnitude for each observation in SS1 by correcting for phase and distance. The same procedure is conducted for TESSC1. To combine the two data sets for each object, we compute the mean of the magnitudes in

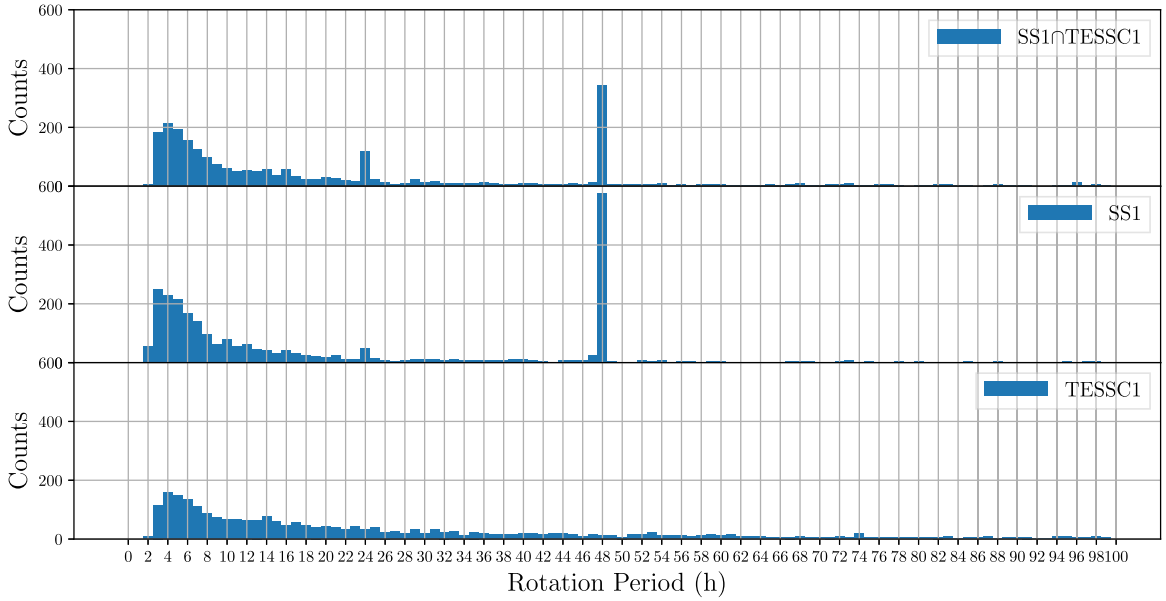

Figure 2. The period distribution for the $n = 3168$ objects in SS1∩TESSC1. The bottom to top panels show the period distribution using only the observations in TESSC1, the period distribution using only the observations in SS1, and the period distribution when combining the observational records of both data sets, respectively. Although we search for solutions up to 10,000 hr, we only show solutions up to 100 hr because there are so few solutions >100 hr.

TESSC1 and then use this to offset the corresponding observational record in the SS1 data set. We elected to normalize to TESSC1 because there are more observations on average in the TESSC1 data set than in the SS1 data set. The photometric errors are unchanged when combining the magnitudes from both data sets. The observational records for each object often contain several apparitions. The apparitions vary between objects, where SS1 and TESSC1 apparitions can be temporally disjoint, or there can be overlap where SS1 observations occur during the same apparition as the TESSC1 observations. Several examples of unphased photometry will illustrate this in Section 5.4. Figure 1(c) plots the distribution of objects in the combined SS1∩TESSC1 data set where the mean number of observations per object is $\mu = 458.9$.

As we will show later in Section 4.2.1, we also examine σ clipping the SS1∩TESSC1 data set at the 2σ and 3σ levels. The σ-clipped variants of this data set contain a total of 1,384,659 and 1,449,961 observations, respectively.

3. Rotation Period Derivation

To derive the rotation periods of asteroids, we use the Lomb–Scargle periodogram (LSP; N. R. Lomb 1976; J. D. Scargle 1982). We employ the GPU implementation outlined in M. Gowanlock et al. (2021), which uses the generalized variant of LSP that uses the floating mean method and considers photometric error when fitting the light curve.

While the periods derived for TESSC1 that are outlined by A. McNeill et al. (2023) also used the generalized LSP algorithm with the photometric error, the rotation periods reported in D. E. Trilling et al. (2023b) ignored the photometric error. Consequently, to ensure consistency between the methods used to derive the rotation periods across all data sets, we rederived the rotation periods for each object in SS1 using the same approach.

Across all data sets, for each object we execute LSP and search a uniformly spaced frequency grid of $n_f = 10^6$ frequencies. For the SS1 and SS1∩TESSC1 data sets, we search in the frequency range $[f_{\min}, f_{\max})$, where $f_{\min} = 0.0048\ \mathrm{day}^{-1}$ and $f_{\max} = 24.0\ \mathrm{day}^{-1}$, respectively. This corresponds to rotation periods of ≈2–10,000 hr, where the rotation period of an asteroid is twice the light-curve period derived by LSP (one light-curve period would only capture half of an asteroid's rotation).

For each of the $n_f = 10^6$ frequencies searched for each object, LSP outputs a power value. In this paper, the frequency with the highest power in the periodogram corresponds to the light-curve period, which is doubled to compute the rotation period. In all that follows, when we refer to the period of an asteroid, we are referring to the *rotation period* and not the light-curve period. This quantity will be reported in hours throughout the paper.

3.1. Period Distribution of SS1∩TESSC1

Using the LSP, we derive the rotation periods of all $n = 3168$ asteroids in SS1∩TESSC1 by selecting the period corresponding to the frequency with the greatest power in the periodogram.

Figure 2 shows the distribution of rotation periods, and we observe that the period distribution for SS1 (middle panel) has a large spike at 48 hr, which is clearly an alias generated by the diurnal signature described in Section 1. Contrasting this with TESSC1 (bottom panel), we observe that there are no spikes at typical aliases, which is to be expected, as the space-based observatory is not impacted by a diurnal observing schedule. When we combine the observational records together to obtain SS1∩TESSC1, we find that there is an overabundance of objects having a 48 hr period and a smaller overabundance at 24 hr. This is surprising, as we would expect that TESSC1 may eliminate a large fraction of period solutions that are derived at the abovementioned aliases. While the number of aliases has decreased when comparing the middle panel to the top panel of Figure 2, the overabundance of period solutions at aliases in SS1∩TESSC1 clearly indicates that the ground-based ZTF data (SS1) negatively impact the derived rotation period solutions.

4. Comparison to LCDB and Improving Derived Period Fidelity

In this section, we compare the fidelity of our period solutions to those in the literature. We begin by comparing the

Table 2
The Fraction of Objects in Three Data Sets, TESSC1, SS1, and SS1∩TESSC1, That Match the Period Solutions in LCDB

Data Set Name	n	Match Fraction (Exact)	Number (Exact)	Match Fraction (Incl. Aliased)	Number (Incl. Aliased)
SS1	222	0.599	133	0.608	135
TESSC1	222	0.739	164	0.806	179
SS1∩TESSC1	222	0.784	174	0.829	184
Upper bound	222	0.820	182	0.878	195

Note. There are $n = 222$ objects in LCDB∩SS1∩TESSC1. The upper bound refers to whether one of the TESSC1, SS1, or SS1∩TESSC1 data sets contains the correct period solution. We report exact matches to LCDB and aliased matches that are within a factor $0.5\times \pm 3\%$ or $2\times \pm 3\%$ of the LCDB period solution.

period solutions in the literature with those derived from the SS1∩TESSC1 data set, which is our baseline approach (Section 3.1). Then, we apply several optimizations to improve the confidence in our period solutions and determine to what extent we can improve derived period fidelity over the baseline.

4.1. Baseline Approach

The LCDB contains asteroid rotation periods reported in the literature and a confidence level (U). In this paper we only consider objects with a quality score of $U = 3$, which are unambiguous rotation period solutions where there is full sampling of the light curve (B. D. Warner et al. 2009). Taking the intersection of the $n = 3168$ objects in SS1∩TESSC1 with the objects in LCDB, we are left with a total of 222 objects. This sample of 222 objects in LCDB is used to determine whether we obtain correct rotation periods, and we refer to this sample of objects as LCDB∩SS1∩TESSC1.

We define an *exact match* to LCDB when an object has a rotation period within 3% of the LCDB period. Often there will be several power spikes in the periodogram where two or more spikes will yield nearly identical power values. Often these occur at periods that are a factor of half or double the true rotation period defined by LCDB; therefore, we also report when our periods are a factor of $0.5 \times \pm 3\%$ or $2 \times \pm 3\%$ of the LCDB period solution. We refer to these as *aliased matches*, although they should not be confused with aliases that are a function of the diurnal observing schedule of a ground-based telescope (Section 1).

Table 2 reports the fraction of matches to LCDB for SS1, TESSC1, and SS1∩TESSC1. We find that the ZTF data set (SS1) has an exact match percentage of ≈60%, whereas TESSC1 has an exact match percentage of ≈74%. This is to be expected, as there are more observations per object in TESSC1 compared to SS1 (Figures 1(a) and (b)), which helps improve the overall match fraction. Furthermore, as described in Section 1, the TESSC1 data set is not susceptible to aliases that are common in SS1. We find that combining SS1 with TESSC1 (SS1∩TESSC1) yields the greatest exact match percentage of ≈78%.

As described in Section 3, we rederived the rotation periods for the SS1 data set outlined in D. E. Trilling et al. (2023b) to include the photometric error and use the generalized variant of LSP. We found a ≈60% match percentage to LCDB. This match percentage is consistent with D. E. Trilling et al. (2023b), which found that for an LCDB quality code $U = 3-$ or better they achieve a match percentage with LCDB of 67%, where they define an exact match to LCDB as a derived rotation period being within 10% of the LCDB rotation period. Given that we use a 3% threshold in this paper and that the samples are different (we only examine those objects in SS1∩TESSC1), our 60% match percentage is roughly consistent with the 67% match percentage reported by D. E. Trilling et al. (2023b).

Table 2 also reports the upper-bound match fraction. We define the upper bound as whether one of the three data sets (SS1, TESSC1, or SS1∩TESSC1) obtains the correct period solution for an object. We find that the upper-bound exact match fraction is 82%, implying that for 18% of the objects none of the three data sets are able to derive the correct period solution. As we will show in future sections, the 18% of objects for which we do not derive the correct period solution are largely low-amplitude objects.

Regarding exact matches, our combined SS1∩TESSC1 data set is within ≈4.4% of the upper bound, and when we include the aliased period solutions that are a factor $0.5\times \pm 3\%$ or $2\times \pm 3\%$ of the LCDB period solution, our results are within ≈5.6% of the upper bound.

Figure 3 shows histograms limited to objects with the exact correct rotation period as summarized in Table 2 as a function of rotation period. By plotting the histogram as a function of rotation period, we can observe whether any of the data sets yield period solutions that are deficient at a particular period range. Plotted in the top panel is the true period distribution from LCDB, and this histogram is plotted behind the three other histograms in the bottom three panels. This makes it straightforward to observe at what rotation periods the three data sets are deficient at finding the correct rotation period, which is denoted by an excess orange outlined bar compared to the blue bar at a given period.

Overall, we find that none of the data sets are able to detect the periods in the smallest bin at 2 hr and that the data sets are generally unable to recover many of the periods at $2\text{ hr} < p < 4\text{ hr}$. This is consistent with A. McNeill et al. (2023), who found that TESS rotation periods were unreliable at ⩽3 hr, and D. E. Trilling et al. (2023b), who found a similar result. Deriving the periods for fast-rotating asteroids is difficult with TESS because of the 30-minute cadence of TESS full-frame images. In the case of ZTF, sparse photometry makes it challenging to constrain asteroid rotation periods when they are rapidly rotating, and it is for this reason that D. E. Trilling et al. (2023b) did not search for rotation periods <2 hr. At rotation periods ≳4 hr, we find no discernible trends; however, because there are only 222 objects in the sample, it is possible that there are trends beyond ≳4 hr, but these cannot be recovered owing to small number statistics.

4.2. Improving Confidence in Period Solutions

As shown in Section 4.1, the baseline approach using SS1∩TESSC1 achieves a match percentage of 78.4%–82.9%. We highlight that obtaining a period match rate of ∼80% is outstanding; therefore, improving the match rate much beyond this level will only recover periods at the margins in a few cases. Thus, in what follows, we attempt to recover periods in these marginal cases, while simultaneously improving the match fraction such that we have greater confidence in the derived period solutions. The results are summarized in Table 3.

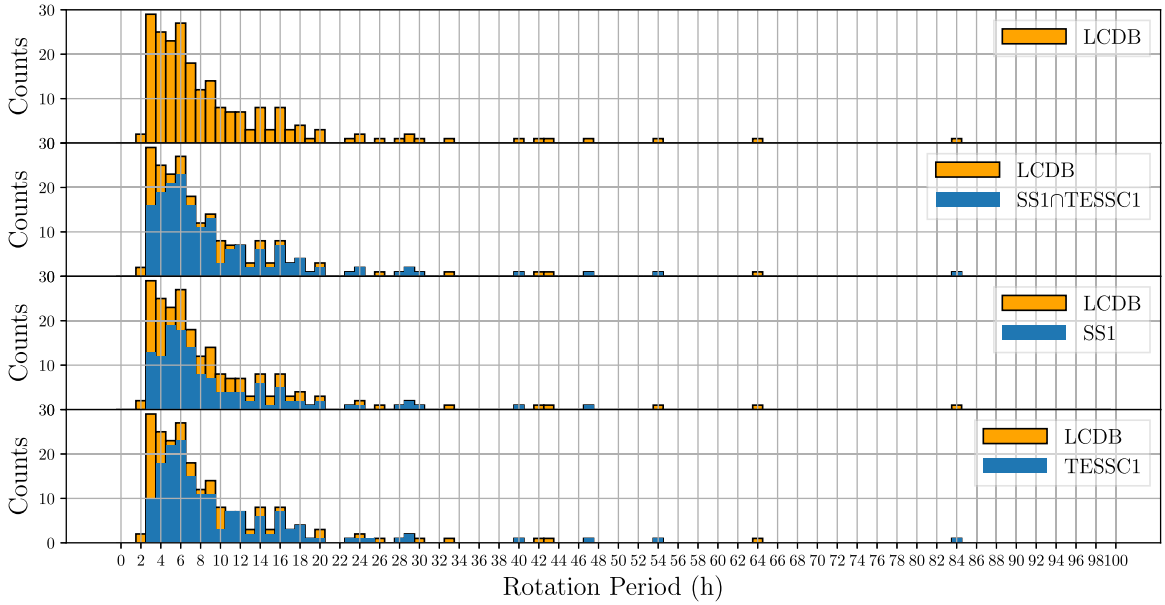

Figure 3. The distribution of correct rotation periods as a function of period for the 222 objects in LCDB∩SS1∩TESSC1 across the four data sets is shown. Each rotation period bin is 1 hr, and the bins are centered on integer period values. The top panel shows the LCDB period distribution, where each object has been assigned an unambiguously correct rotation period. The three lower panels from top to bottom are as follows: SS1∩TESSC1, SS1, and TESSC1, where the (correct) LCDB distribution is plotted behind the period distribution shown for each of these three data sets. Only those objects whose periods exactly match the LCDB period are plotted.

Table 3
Comparison of the Fraction and Number of Matches of the Objects in LCDB∩SS1∩TESSC1 as a Function of Optimizations (1)–(7)

Method	Sections	$n_{\rm LCDB}$	$p_{\rm min}$	$n_{\rm amp}$	Exact Match Frac.	Exact Number	Incl. Aliased Match Frac.	Incl. Aliased Number	Utility
Baseline: SS1∩TESSC1 (Table 2)		222	⩾2 hr	222	0.784	174	0.829	184	
(1) 2σ clipping	4.2.1	222	⩾2 hr	222	0.797	177	0.842	187	✓
(2) 3σ clipping	4.2.1	222	⩾2 hr	222	0.788	175	0.833	185	✗
(3) Periodogram masking: aliases (0.1 hr mask width)	4.2.2	222	⩾2 hr	222	0.788	175	0.833	185	✓
(4) Periodogram masking: $p < 0.9(t_{\rm window})$ hr	4.2.3	222	⩾2 hr	222	0.784	174	0.829	184	✗
(5) Replace with TESSC1 period at high confidence level	4.2.4	222	⩾2 hr	222	0.793	176	0.838	186	✓
(6) Replace with SS1 period at high confidence level	4.2.5	222	⩾2 hr	222	0.779	173	0.824	183	✗
(7) Excluding amplitudes ⩽0.075	4.2.6	222	⩾2 hr	203	0.793	161	0.823	167	✓
{1, 3, 5, 7} Combined	4.2.7	222	⩾2 hr	196	0.827	162	0.852	167	✗
{1, 3, 5, 7} Combined	4.2.7	206	⩾3 hr	182	0.868	158	0.890	162	✓
{1, 3, 5, 7} Combined	4.2.7	173	⩾4 hr	151	0.881	133	0.907	137	✗

Note. $n_{\rm LCDB}$ refers to the number of objects in LCDB∩SS1∩TESSC1 after discarding objects with periods below the period cutoff ($p_{\rm min}$), and $n_{\rm amp}$ refers to the number of objects in the sample after discarding those with light-curve amplitudes below the threshold using method (7). When $n_{\rm LCDB}$ and $n_{\rm amp}$ are equivalent, this implies that no amplitude threshold was utilized. The match fraction is the total number of objects with a correct rotation period as a ratio of $n_{\rm amp}$. The utility column refers to whether the optimization should be used (✓) or not (✗).

4.2.1. Sigma Clipping

The SS1 data set reported in D. E. Trilling et al. (2023b) has been filtered upstream from the SNAPS broker by the ANTARES broker (A. Saha et al. 2016; T. Matheson et al. 2021). In this process, observations with low real-bogus scores are removed from consideration. Therefore, the observational records for each object are largely free from contamination by poor observations. The TESSC1 data set reported in A. McNeill et al. (2023) has been σ clipped to remove outliers in the observational records for each object.

While each of the SS1 and TESSC1 data sets has individually had outliers removed from the observational records for each object, when combining the data sets together to create SS1∩TESSC1, new outliers in the combined observational record may be introduced. Consequently, we investigate removing outliers by removing observations with magnitudes that exceed 2σ or 3σ from the median magnitude. We do this using the σ-clip function in Astropy (T. P. Robitaille et al. 2013), `astropy.stats.sigma_clip`, with a single iteration.

We compare the match fraction to LCDB for both 2σ and 3σ levels. Table 3 shows that 3σ clipping (the standard level) is able to recover the correct period for an additional object over the baseline for both exact and aliased period solutions. Furthermore, 2σ is able to recover the correct period for three additional objects over the baseline for both exact and aliased period solutions. This demonstrates that new outliers were introduced when combining the data sets and that σ clipping

can improve the total match fraction. We caution that σ clipping below the 2σ level will remove observations from objects that have a high light-curve amplitude, so we do not investigate removing outliers below the 2σ level.

4.2.2. Periodogram Masking: Aliases

The diurnal observing schedule of ground-based telescopes imparts a periodic signal into the periodogram. Excluding periods at aliases such as $p \in \{16, 24, 48\}$ hr is straightforward—if the peak power in the periodogram is at one of these aliases, then a secondary peak is selected. The drawback of this method is that a true period at these aliases is rejected despite the fact that asteroids may truly have these rotation periods (e.g., see Figure 2, bottom panel or Figure 3, top panel).

Recently, N. Erasmus et al. (2021) used masking to detect slowly rotating asteroids. M. W. Coughlin et al. (2021) used the masking method when examining variable stars in ZTF DR2.[5] D. Kramer et al. (2023b) compared several methods that remove aliases and found that masking is both the simplest and most effective method on the Legacy Survey of Space and Time (LSST) Solar System Products Data Base (SSPDB; M. Juric et al. 2021), which is used to simulate the LSST observational records and was competitive with the Monte Carlo method on the SS1 data set. Consequently, we examine the use of the masking method here.

Figure 4 shows the match fraction to LCDB as a function of the mask width around the {16, 24, 48} hr aliases. Here the mask widths are centered on the abovementioned aliases and are reported as ± the mask width as shown on the horizontal axis. We find that masking does not significantly improve the overall match fraction. In addition, as the mask width increases, the total match fraction decreases. This is expected, as we increasingly eliminate larger fractions of the frequency range that may contain the true period of an object.

This result is consistent with that of D. Kramer et al. (2023b), which showed that masking only improved the match percentage from 64.8% to 65.8% on the SS1 data set. The reason this occurs is that the secondary peak selected outside of the masked region of the periodogram is often incorrect. Despite this result on the SS1 data set, D. Kramer et al. (2023b) showed that on the SSPDB data set masking improved the match percentage from 57.9% to 74.5%, demonstrating that the method is expected to work well for the forthcoming LSST catalog.

We reiterate that masking is irrelevant for the TESSC1 data set, as it does not suffer from a diurnal observing schedule, so when we combine the SS1 and TESSC1 data sets, TESSC1 may eliminate aliases, which diminishes the utility of the masking method. Despite this, we expect that recovering the periods for a few additional (marginal) objects is worthwhile even if the number of objects is not statistically significant here. Thus, when we combine optimizations, we use this method and apply masks at {16, 24, 48} ± 0.1 hr.

4.2.3. Periodogram Masking: $p \geqslant 0.9(t_{window})$

Ground-based observatories, such as ZTF and the forthcoming Rubin Observatory, will have a long observational baseline for each object. This will allow deriving long rotation periods for a small fraction of asteroids. In contrast, TESSC1 has a much shorter observational baseline for each object. The telescope stares at a sector and then transitions to the next sector, where it is unlikely that an asteroid is captured in more than one sector. As a result, the maximum rotation period that can be detected with TESSC1 is roughly 816 hr (see A. McNeill et al. 2023 for more information). Let $t_{window} = t_{end} - t_{start}$, where t_{start} and t_{end} refer to the first and last times that the object has been observed, respectively. Then, if $p > t_{window}$ for an object, this implies that the light curve will be poorly sampled, as only a partial light curve will be produced. Consequently, we mask all periods for an object where $p \geqslant 0.9(t_{window})$ such that we reject periods that are within 10% of the observing window, as periods derived within the last 10% of the observing window are well-known to be unreliable (A. McNeill et al. 2023).

From Table 3 we find that this method has no impact on the total match fraction. In this set of 222 objects, there were no instances where $p \geqslant 0.9(t_{window})$. While this filter was important for deriving asteroid periods for TESSC1, when the data record is combined with SS1, each object is assigned a much larger t_{window}. Thus, this method is unnecessary when applied to the SS1∩TESSC1 data set.

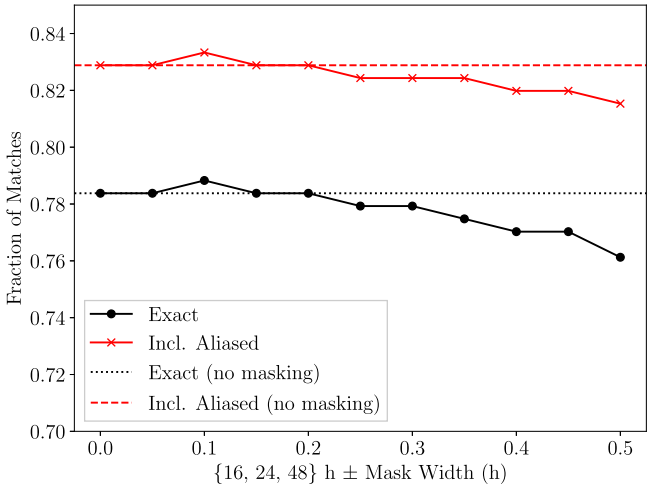

Figure 4. The fraction of exact and aliased matches is shown as a function of mask width. Masks are applied at $p \in \{16, 24, 48\}$ hr. The mask ranges are reported as ± the mask width on the horizontal axis. For example, a mask width of 0.1 around the 16 hr alias refers to masking the range 15.9–16.1 hr. The mask width of 0 refers to the match fraction without any masking. For clarity, the exact and aliased match fractions without masking are shown as the black dotted and red dashed horizontal lines, respectively.

4.2.4. Period Replacement with High-confidence TESSC1 Solutions

A. McNeill et al. (2023) showed that objects having a sufficient number of observations and LSP power are more likely to produce the correct period solution than objects having few observations and low LSP power. This is intuitive, as the greater the number of observations, the more likely that a good model fit is obtained, and furthermore, greater LSP power typically indicates that the periodogram has a high signal-to-noise ratio and that the light curve has a reasonably high amplitude. A. McNeill et al. (2023) showed that, without assigning a threshold for the LSP power, they achieved a match fraction to LCDB of 0.65. In contrast, for those objects observed ⩾200 times and having an LSP power ⩾0.2, the match fraction increases to 0.85, and this further increases to 0.91 if aliased periods are also considered.

Given the above, if a period in the TESSC1 data set has a high confidence of being correct (⩾200 observations and an LSP power ⩾0.2), then we simply select this period instead of

[5] https://www.ztf.caltech.edu/ztf-public-releases.html

using the period derived from SS1∩TESSC1. This approach also has the benefit of potentially validating solutions at the {16, 24, 48} hr aliases. From Table 3 we find that this method is able to recover two additional objects relative to the baseline.

4.2.5. Period Replacement with High-confidence SS1 Solutions

Similarly to period replacement with high-confidence TESSC1 solutions, we also examine replacement using SS1. Figure 5 shows a similar heatmap to that outlined in A. McNeill et al. (2023), which plots a cumulative heatmap of period matches to LCDB as a function of the observation threshold and LSP power threshold. With unconstrained observational and LSP power thresholds we obtain an exact match fraction of 0.597, but this increases to 0.85 when there are $\gtrsim$100 observations and an LSP power of 0.7.

If a period in the SS1 data set has a high confidence of being correct ($\geqslant$105 observations and an LSP power $\geqslant$0.7), then we select this period instead of using the period derived from SS1∩TESSC1. From Table 3 we find that this method does not recover any additional correct rotation periods relative to the baseline. Furthermore, based on Figure 5(b), this only captures a small fraction of the objects in the sample because the number of observations and number of LSP power thresholds need to be very high, which limits its utility for period replacement. Consequently, when we combine optimizations later, we do not use this method.

Note that future SNAPS data releases will contain significantly more ZTF observations per object, which will likely increase the utility of this approach.

4.2.6. Excluding Low-amplitude Objects

We now examine the light-curve amplitude distribution of the objects in the data sets. Amplitudes are derived using the sinusoidal fit of the light curve for each object using the peak and trough of the curve. On average, this will underestimate the amplitudes of asteroids, as a sinusoid will often not capture the peak and trough within a light curve. However, this is preferable to alternative methods, such as taking the difference between the minimum and maximum magnitudes in a light curve, which may dramatically overestimate the amplitude owing to outlying measurements with low and/or high magnitudes.

Figure 6(a) plots the amplitude distribution of the 3168 objects in the SS1∩TESSC1 data set, which shows that there is a deficit of amplitudes $\leqslant$0.1 mag. This is due to bias in period finding, where asteroids with a spherical morphology will have a light-curve amplitude of $\sim$0 mag, and thus it is not possible to derive rotation periods for these objects. Furthermore, light curves that are fit with a poor period may produce a low signal-to-noise periodogram, which often generates a relatively flat (low-amplitude) fit to the time series. In addition, asteroids with a pole-on orientation of any shape can produce a flat light curve. Consequently, the deficit of amplitudes $\leqslant$0.1 mag is largely unphysical, and it is difficult to know the true distribution of objects with amplitudes $\leqslant$0.1 mag.

Comparing Figure 6(b) on the LCDB∩SS1∩TESSC1 data set to Figure 6(a), we find that there are a greater fraction of amplitudes $\leqslant$0.1 mag. This is expected, as the LCDB∩SS1∩TESSC1 sample contains objects vetted from LCDB, and hence the sample should be a closer match to the real amplitude distribution, which cannot be determined owing to the abovementioned presence of asteroids that are roughly spherical in shape. Figure 6(c) shows the amplitude distribution of the LCDB∩SS1∩TESSC1 data set that is limited only to the 174 objects that are an exact match to LCDB. Because all objects have been assigned the correct period, we plot this as a sanity check to demonstrate that a similar deficit of amplitudes $\leqslant$0.1 mag is present as shown in Figures 6(a)–(b). If this deficit was not shown in this plot, then this would indicate that spherical objects and poor light-curve fits are not the cause of the deficit described above.

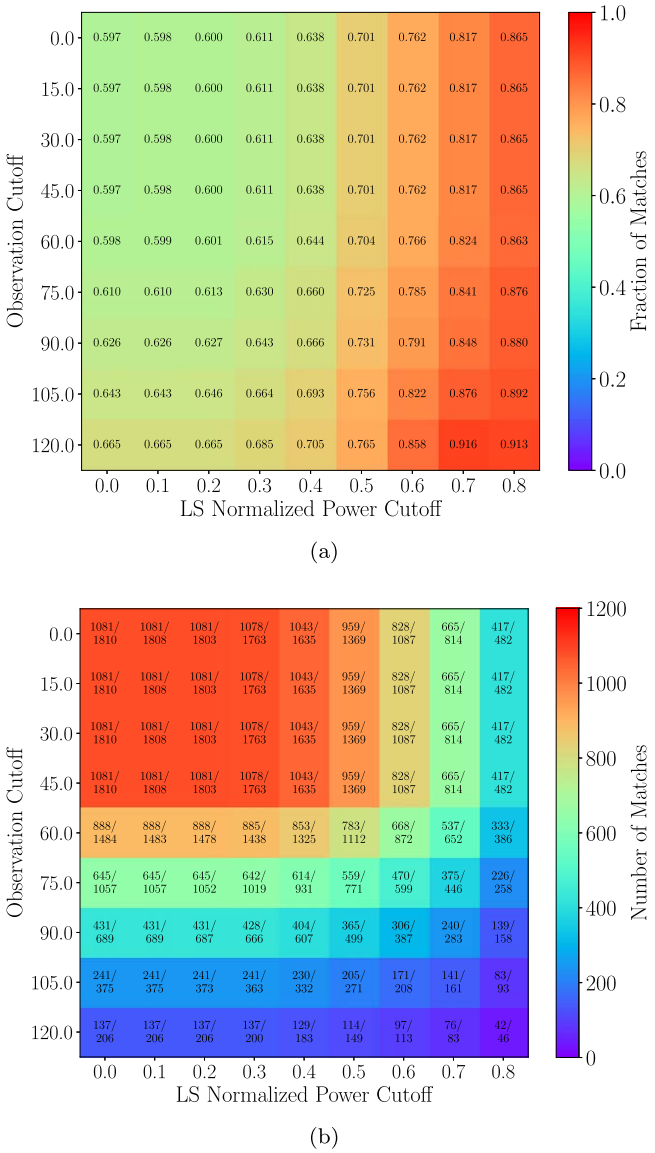

Figure 5. Cumulative heatmaps showing the rotation period match fraction (and number) of objects in SS1 as a function of both the number of observations and the LSP power. In the sample, there are a total of 1810 objects, which are those that are found in LCDB ($U = 3$) and SS1. (a) The exact rotation period match; (b) the number of objects in the bins shown in panel (a).

An additional challenge when fitting light curves is photometric error. We find that many of the observations in the TESSC1 data set have a photometric error >0.1 mag (the photometric error with SS1 is lower), and so if we assume that the average error is roughly 0.1 mag, then we would not expect to be sensitive to objects with amplitudes $\leqslant$0.1 mag.

Figure 7 shows two plausible scenarios for the amplitude distribution at <0.1 mag. We describe two hypotheses as follows:

1. *"Sphere-Abundant Hypothesis":* The first scenario (red dashed curve) shows the case where the amplitude histogram

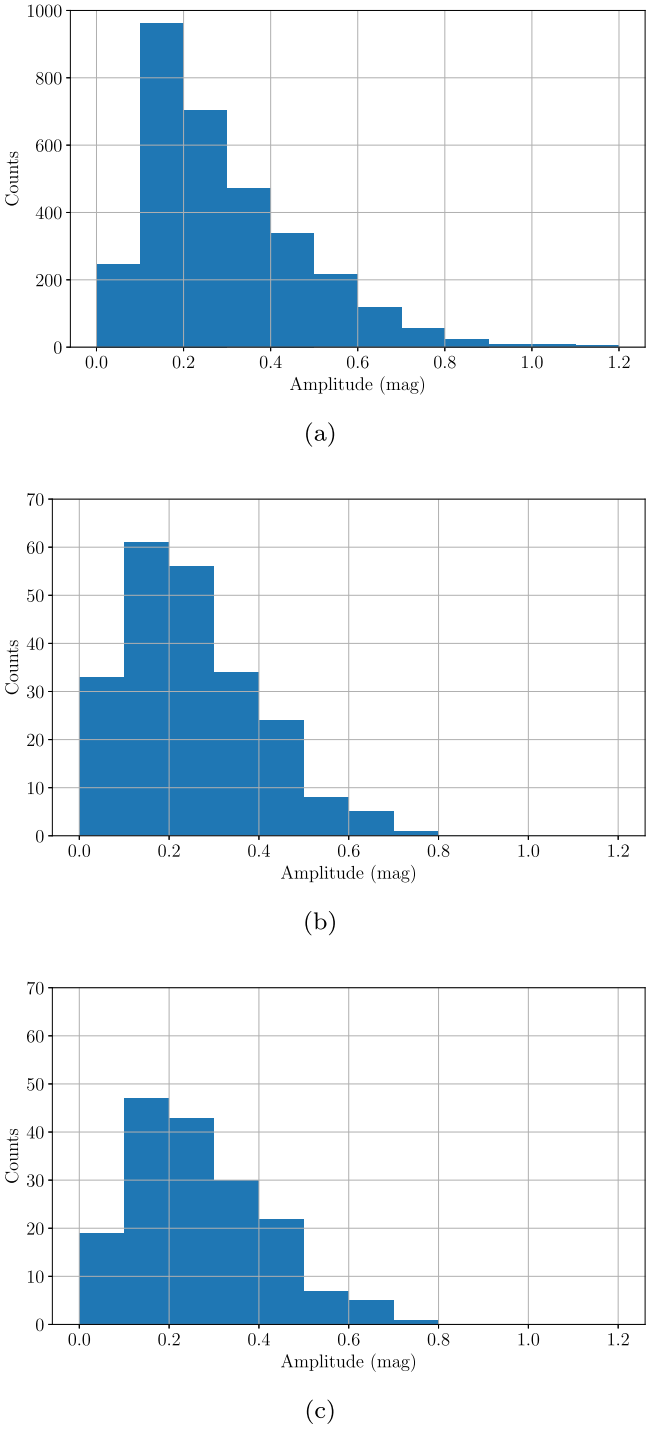

Figure 6. Light-curve amplitude distributions in 0.1 mag bins. (a) Amplitudes of the 3168 objects in the SS1∩TESSC1 data set. (b) Amplitudes of the 222 objects in the LCDB∩SS1∩TESSC1 data set. (c) Same as panel (b), except limited to the 174 exact matches to the LCDB, and so panels (b) and (c) are plotted on the same scale for comparison purposes.

should increase monotonically with decreasing amplitude. If this hypothesis is correct, this implies that there are numerous asteroids that are spherical in morphology and that the amplitude distribution will be difficult to accurately constrain.

2. *"Sphere-Limited Hypothesis":* The second scenario (black dotted curve) is where there are fewer objects with amplitudes <0.1 mag than at 0.1–0.2 mag, implying that the vast majority of asteroids are likely to be at least slightly elongated.[6] If this hypothesis is correct, then the amplitude distributions with a deficit at <0.1 mag shown in Figure 6 are indicative of the true distribution.

In the context of the above, we describe implications for rotation period derivation. If we assume that the "Sphere-Abundant Hypothesis" is correct, then this implies that we are unable to correctly assign ∼25% of the objects their correct amplitude; thus, it is likely that the derived rotation period for a significant fraction of these objects is also incorrect. One important caveat to note is that it is possible to derive the correct rotation period for an asteroid but still derive a poor amplitude (i.e., asteroid light curves may not always be best fit by LSP sinusoids); therefore, we would not expect that 25% of the rotation periods are incorrect based on an analysis of light-curve amplitudes alone. If the "Sphere-Limited Hypothesis" is correct, then this implies that light-curve amplitudes cannot be the dominant reason that we obtain incorrect period solutions.

Recall that the baseline finds that there is a 0.784 exact match fraction of rotation periods (Table 3). Thus, up to ≈25% of rotation periods may be incorrect, which is consistent with the fraction of asteroids missing from the distribution having an amplitude <0.1 mag (assuming that the "Sphere-Abundant Hypothesis" is correct). We conclude that a substantial fraction of asteroids that have an unreliable derived rotation period may be due to spherical asteroids with rotation periods that are intractable to constrain.

Figure 8(a) plots the fraction of objects that match LCDB as a function of the amplitude cutoff. Unlike the other optimizations presented in this section, by making a cut on the amplitude, we lose some number of objects that have been assigned a correct rotation period, and so we also plot the number of objects recovered in Figure 8(b). Consequently, we aim to select an amplitude cutoff that reaches a trade-off of yielding a high match fraction without eliminating too many objects from the sample with a correct rotation period. For instance, in an extreme case, we could achieve a ∼90% match fraction by removing objects having an amplitude <0.3 mag, but we would only recover 65 asteroids.

Figure 8(a) reveals an interesting trend regarding exact versus aliased matches where we find that the match fractions for both exact and aliased matches converge when we exclude objects with amplitudes <0.175 mag. This implies that the LSP period-finding algorithm produces an ambiguous solution (within a factor $0.5\times \pm 3\%$ or $2\times \pm 3\%$ of the LCDB period solution) when there is a low amplitude, and this is to be expected since low-amplitude period solutions also produce low-power periodograms, which implies that the peak power in the periodogram is not significantly differentiated from powers at other candidate frequencies. But when the amplitude is high, the ambiguity seems to be reduced or eliminated.

Based on the above, we select an amplitude threshold that removes all objects from the sample at <0.075 mag, which allows for some of the objects in the 0–0.1 mag bin to be included in the sample. This reaches a trade-off between maximizing the match fraction and maximizing the number of objects that have been assigned a correct rotation period. With this amplitude threshold, we find that there are a total of 203 objects (thus 19 objects are excluded), and this yields an exact (aliased) match fraction of 0.793 (0.823) corresponding to 161 (167) objects, respectively.

[6] While not directly applicable to main-belt asteroids, this hypothesis was proposed for trans-Neptunian objects (P. H. Bernardinelli et al. 2023).

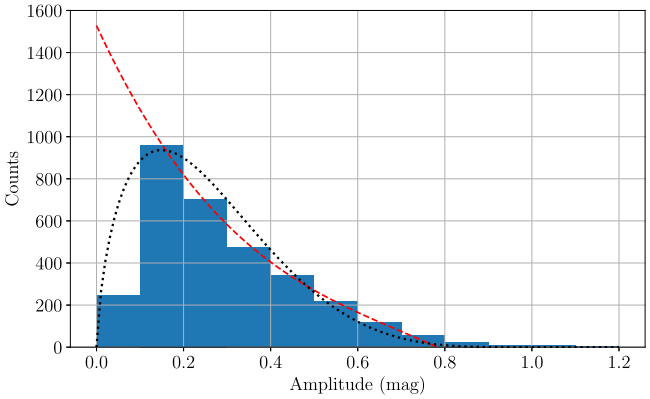

Figure 7. The same histogram as in Figure 6(a), except where two example models are shown that hypothesize the distribution of asteroids with amplitudes <0.1 mag. The red dashed curve shows the scenario where spherical-like objects are the most abundant. Using a β-distribution, the black dotted curve shows the scenario where there are fewer objects at <0.1 mag than at 0.1–0.2 mag.

4.2.7. Summary: Combining Optimizations and Quantifying Period Confidence

The methods described above are used to recover the rotation periods of asteroids at the margins where simply combining the SS1 and TESSC1 data sets are unable to produce the correct period. If we select and combine the methods that were deemed to have utility above, which are (1), (3), (5), and (7) in Table 3, then we obtain an exact match fraction of 0.827, which increases to 0.852 when aliased solutions are considered correct.

Recall that the minimum rotation period search range is $p_{\min} = 2$ hr. However, A. McNeill et al. (2023) demonstrated that periods derived with TESSC1 are unreliable at $p < 3$ hr. Consequently, we examine the fraction of correctly derived rotation periods when we exclude objects in LCDB having $p < 3$ hr, which yields a total of 206 objects in the sample. Of those 206 objects, 182 had a light-curve amplitude >0.075 mag. We find that the exact and aliased match fractions are 0.868 and 0.890, respectively, which corresponds to recovering the rotation periods of 158 and 162 asteroids. Thus, we are able to further increase the match fraction using $p_{\min} = 3$ hr compared to $p_{\min} = 2$ hr. We also examined setting $p_{\min} = 4$ hr, but this yields a dramatic loss in the number of correct asteroid rotation periods recovered (see Table 3).

Based on this analysis, we use the combined optimizations with $p_{\min} = 3$ hr because it yields a trade-off between maximizing the match fraction and maximizing the number of correct periods recovered.

To reiterate, we are able to recover the rotation periods with a 0.890 confidence when aliased solutions are included. However, the optimizations do not completely explain which properties of a light curve may indicate that it has a high probability of being correct. To this end, Figure 9(a) plots the exact match fraction using the set of optimizations and $p_{\min} = 3$ hr, as a function of the number of times an object has been observed and the LSP power, where a higher LSP power typically implies that there is a high signal-to-noise ratio when selecting the maximum power in the periodogram for each object. The heatmap is cumulative, so the upper left corner contains all of the objects that are correct (158) without any cut on the number of observations or LSP power, and the number of objects in each bin decreases as the threshold on the number of observations or LSP power increases.

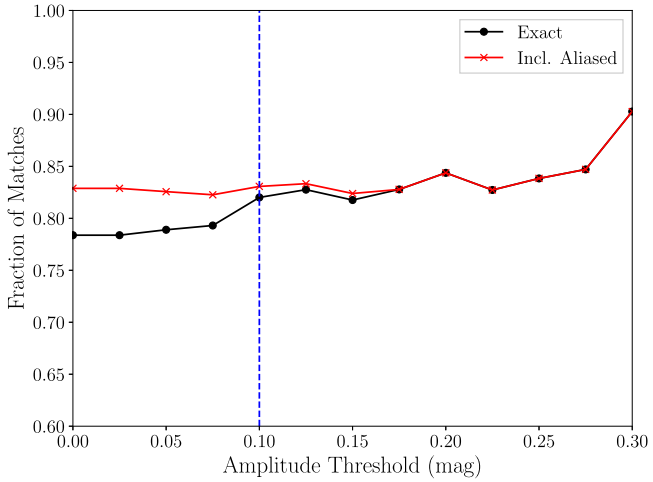

(a)

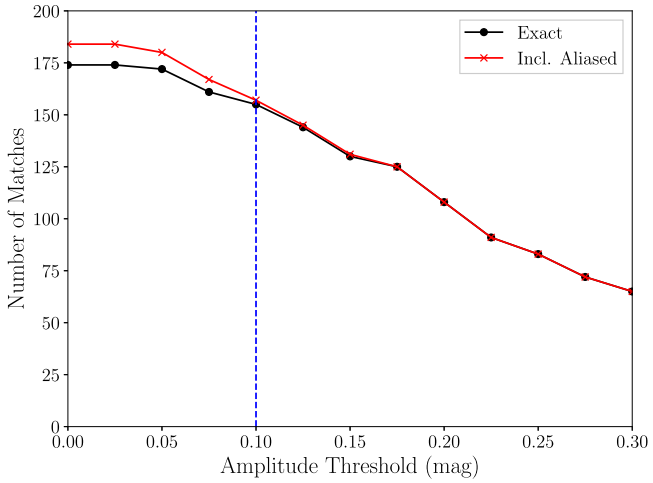

(b)

Figure 8. (a) The fraction of rotation period matches (exact and including aliased) with LCDB as a function of the amplitude threshold. (b) The number of objects corresponding to panel (a). The dashed blue horizontal line demarcates amplitudes at 0.1 mag as described in the text accompanying Figure 6.

We find that we can obtain a 100% match to LCDB when we have at least 375 observations and an LSP power of 0.7. Figure 9(b) shows the number of objects in the bins, indicating that there are 50 objects with correct rotation periods that have been observed at least 375 times and have an LSP power $\geqslant$0.7. While this is a very high observational threshold to obtain a perfect match fraction, a 0.90 match fraction can be obtained using more reasonable constraints: $\geqslant$150 observations with an LSP power $\geqslant$0.3. Figures 9(b) and (d) are the same as Figures 9(a) and (c), except that they include aliased matches as well.

The heatmaps above showed the cumulative match fraction and number of matches for each grid cell where there are fewer objects in the bins as the observation threshold and LSP power increase. To investigate which is more critical—the number of observations or LSP power—and to better understand the impact of these parameters on rotation period discovery, we now examine differential heatmaps, where each grid cell only includes objects with a range of observations and LSP powers that are minimum and maximum values that define each grid cell. Thus, compared to the cumulative heatmaps, the number

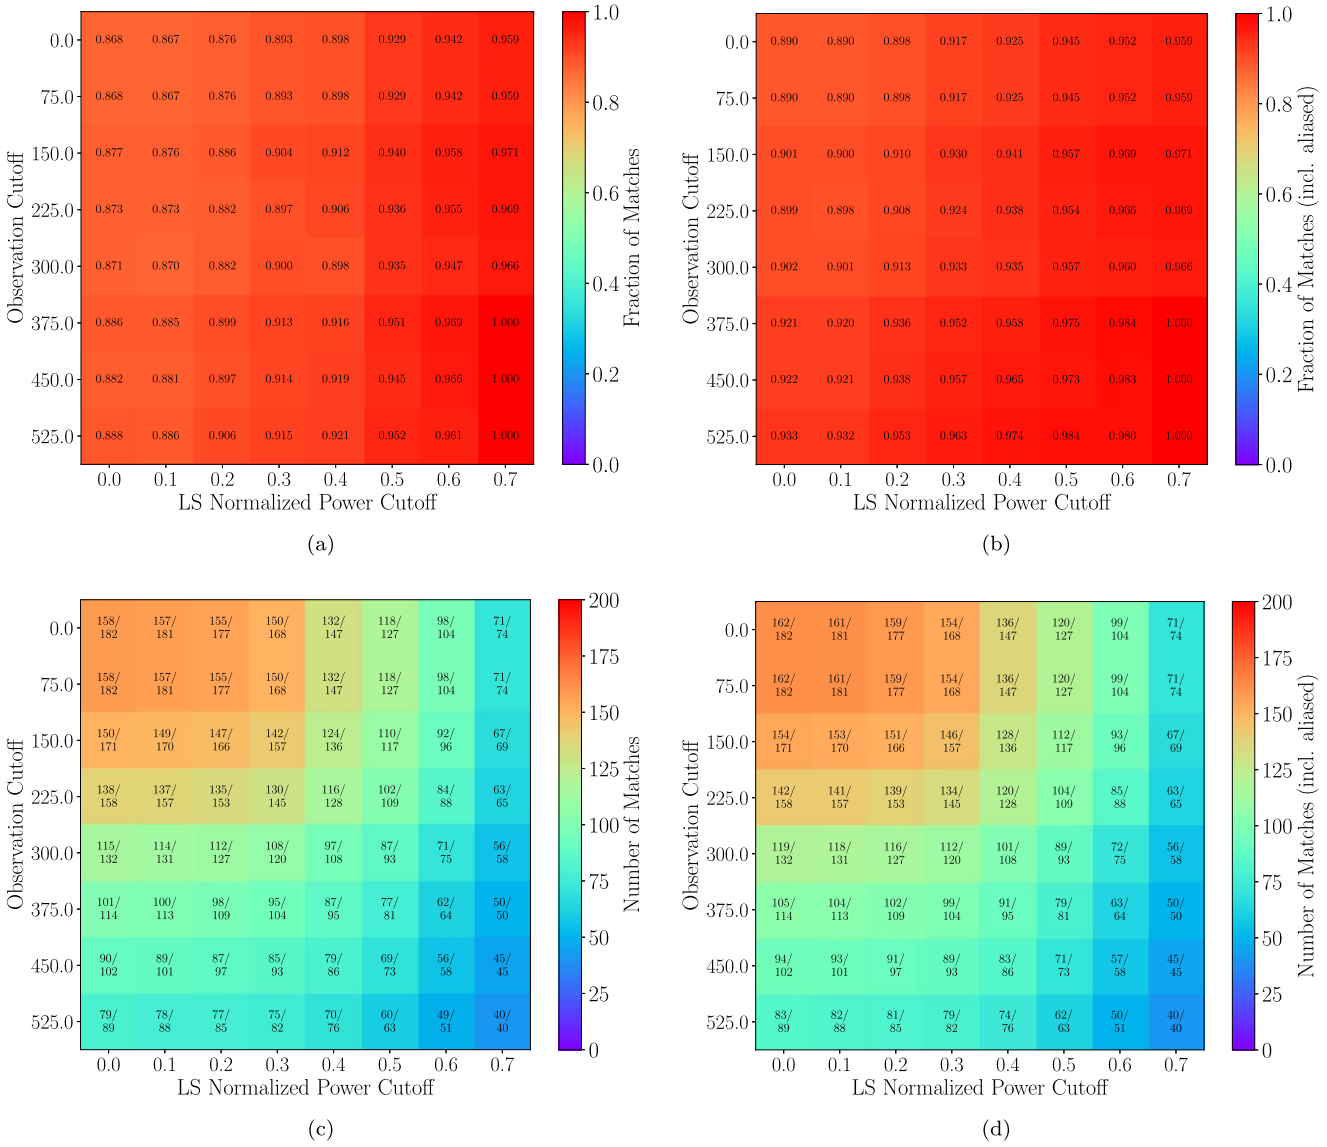

Figure 9. Cumulative heatmaps showing the rotation period match fraction (and number) of objects in LCDB∩SS1∩TESSC1 as a function of both the number of observations and the LSP power. The heatmap uses the data from the combined optimizations outlined in Table 3 with $p_{\min} = 3$ hr. Panels (a) and (c) show the exact rotation period match and number, respectively. Panels (b) and (d) are the same as panels (a) and (c) but include the aliased rotation periods as well.

of objects in each bin will not monotonically decrease as the number of observations and LSP power increases.

Figure 10 shows a differential version of the heatmap shown in Figure 9. This allows us to observe which bins obtain a match fraction below the average across all of the bins (i.e., 0.868 without any observational or LSP power cutoff). There are only 182 objects in the sample, and so the differential heatmap is poorly sampled—thus, to draw any meaningful conclusions, we increased the bin size compared to Figure 9. Considering exact matches, we find that when there are at least 400 observations and an LSP power $\geqslant$0.666 we achieve a perfect match to LCDB. The three bins shown in Figure 10(a) at an LSP power <0.333 are poorly sampled as shown in Figure 10(c), and so while we would expect that the match fraction would increase as we increase the number of observations, we observe the opposite. In this LSP power range, there are only 7, 9, and 0 objects in the bins, respectively, and so small number statistics limits our ability to interpret this result. Lastly, comparing Figures 10(a) and (b), we find that the bin with 800–1200 observations and an LSP power of 0.333–0.666 has a 0.500 match fraction, whereas it has a 0.875 match fraction when aliased matches are included. This shows that the three objects in this bin were assigned a period that was a factor $0.5\times \pm 3\%$ or $2\times \pm 3\%$ of the LCDB period solution, and so the 0.500 match fraction is misleading in Figure 10(a), as objects with $\geqslant$800 observations and an LSP power $\geqslant$0.333 should be considered reliable.

4.3. A Conservative Approach to Biasing against Incorrect Rotation Periods

Sections 4.1–4.2 reported the rotation period match fraction as compared to LCDB. The combination of methods summarized in Table 3 aims to maximize the number of objects with correct periods. However, it does not consider biasing against incorrect rotation periods. We address this drawback by proposing a metric that reports a null result if there is a sufficiently high probability that a period is incorrect.

To include biasing against incorrect rotation periods, we reframe the analysis as a binary classification problem. Each object has three properties that determine whether the period is likely to be correct (light-curve amplitude, LSP power, and

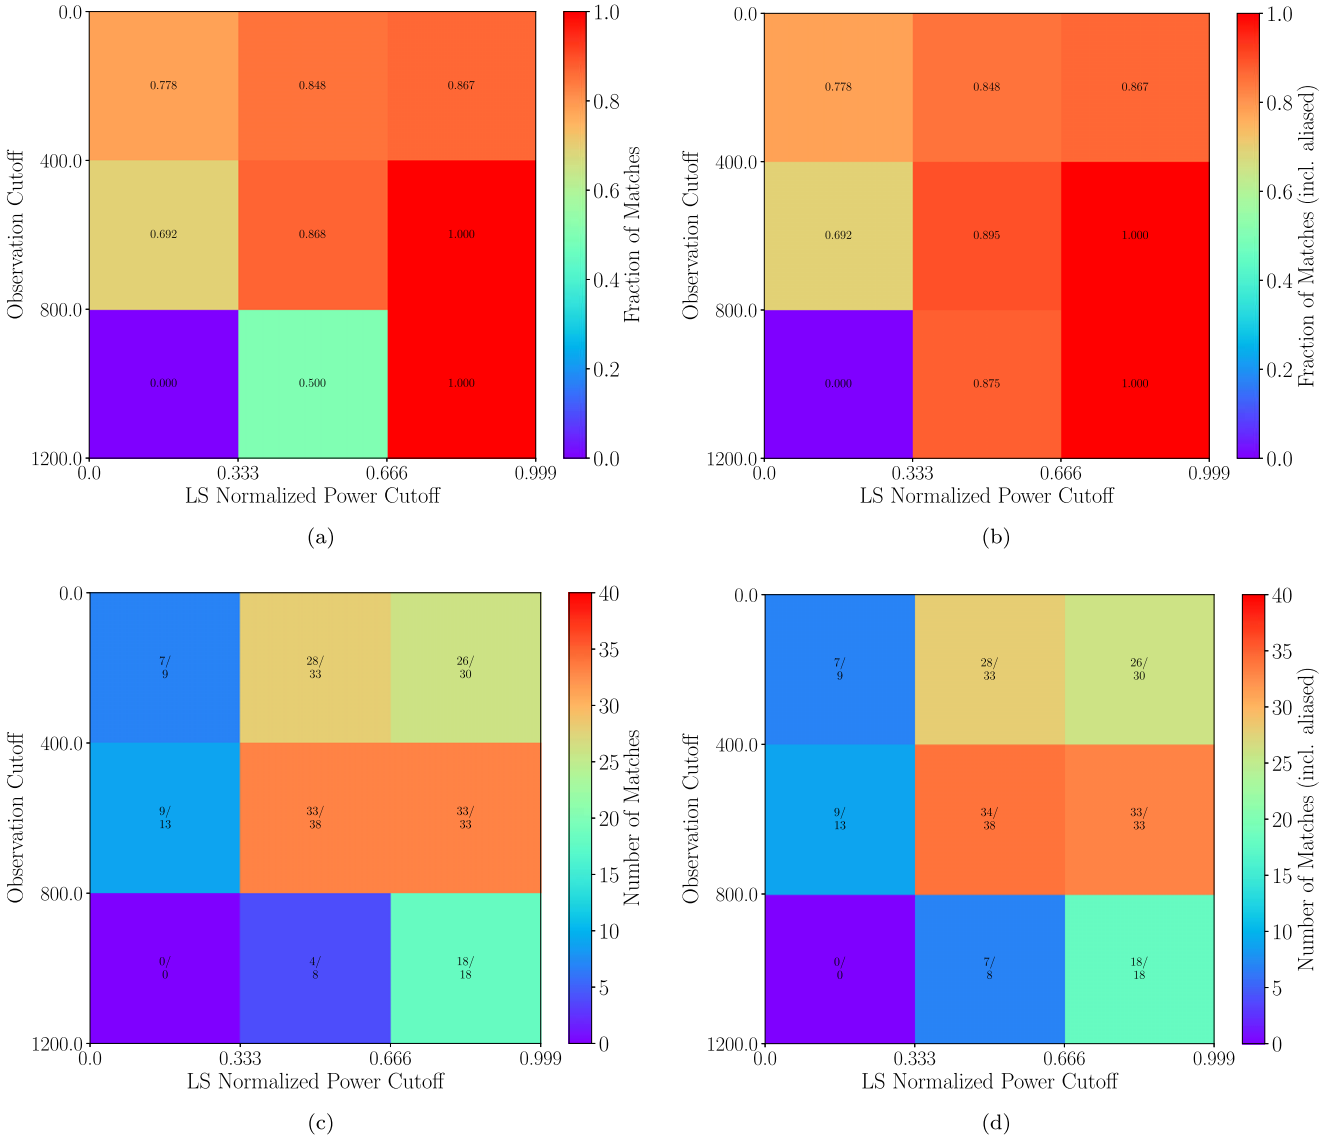

Figure 10. The same as Figure 9, but showing the differential heatmap. Because there are only 182 objects in the sample, the differential heatmap is poorly sampled and is thus very sparse. Consequently, we plot the heatmap using large bins such that we may obtain a significant number of objects in each bin. Caution is warranted, as the number of objects in the bins remains low (e.g., there are no objects with 800 hr $< p <$ 1200 hr and an LSP power 0–0.333).

number of observations), and based on these properties we assign each object a predicted label of either "incorrect period" or "correct period." Then, we compare the period solutions to LCDB to determine whether the labels are correct, where we aim to minimize the number of objects assigned a correct label that are incorrect (a false positive). The metrics described below are those used in supervised machine learning for assessing the quality of a model's ability to predict binary class labels (G. Cabrera-Vives et al. 2017). For clarity, we note that we are not training a machine learning model; rather, we are incorporating metrics from the field of machine learning that are useful for our purposes.

We outline four possibilities regarding the outcome of binary classification as follows:

1. *True Positive (TP):* An object was assigned a correct label, and the rotation period is correct.
2. *False Positive (FP):* An object was assigned a correct label, but the rotation period is incorrect.
3. *True Negative (TN):* An object was assigned an incorrect label, and the rotation period is incorrect.
4. *False Negative (FN):* An object was assigned an incorrect label, but the rotation period is correct.

Using the TP, FP, TN, and FN rates, the following metrics can be computed, where 1.0 is the best (maximum) value for each:

1. Accuracy: (TP+TN)/(TP+FP+TN+FN).
2. Precision: TP/(TP+FP).
3. Recall: TP/(TP+FN).

Observe that the accuracy metric is very similar to the results reported in Sections 4.1–4.2 that examine the match fraction to LCDB, and so we are largely uninterested in this metric here. The recall metric does not consider false positives, and so it is not very useful in this context. The precision metric is the most appropriate, as it biases against reporting false positives. In other words, if we have a precision value close to 1.0, this suggests that the number of false positives (objects assigned a correct label but with an incorrect rotation period) has been minimized.

Observing Figures 8–10, which show the relationship between match fraction and light-curve amplitude, LSP power,

and number of observations, it is clear that perfect precision can be obtained with sufficiently high thresholds for these three properties, as there are several instances that yield a 1.0 match fraction and therefore do not have any false positives.

For an object to be assigned a correct label, it must have a sufficiently high amplitude, number of observations, and LSP power, which are the properties examined in Section 4.2. We use the periods derived by optimizations "{1, 3, 5, 7} Combined" from Table 3 with $p_{\min} = 3$ hr (a total of 182 objects). In the analysis, we only consider exact matches and exclude aliased matches.

We perform a grid search over the parameter space to select thresholds for the following parameters: light-curve amplitude, LSP power, and number of observations. We then use these thresholds to determine whether the object is assigned a predicted correct or incorrect class label. The range of amplitudes searched is [0.075, 0.5] with a step size of 0.025 mag, the range of number of observations searched is [50, 400] with a step size of 10 observations, and the range of LSP power values searched is [0, 0.5] with a step size of 0.05. Each object is assigned a correct label only if it meets all three thresholds.

To summarize, we searched 18 amplitude, 36 observation, and 11 LSP power values, yielding a total of 7128 configurations. Each of these configurations assigns all objects a class label, which are used to derive accuracy, precision, and recall statistics. Figure 11 plots the relationship between accuracy and precision for each of the configurations searched, where we are interested in maximizing the precision since it minimizes false negatives. We make the following observations regarding Figure 11:

1. Higher accuracy is obtained at the expense of lower precision. There are several configurations where perfect (1.0) precision is obtained; however, low accuracy implies that there are many false negatives, which would have the effect of recovering far fewer objects than at higher accuracy thresholds. So perfect precision is not ideal unless one is only interested in correctly deriving the periods of a few objects.
2. There are several configurations that achieve very high precision and reasonably high accuracy, and one of the configurations is denoted by the dotted red lines in the figure. This example configuration has a precision of 0.957 and an accuracy of 0.714, which is achieved with the following thresholds: an amplitude of 0.125 mag, 120 observations, and an LSP power of 0.5. This configuration has a true-positive rate of 111 and an false positive rate of 5.

Using the above thresholds, we confidently assign incorrect and correct labels to objects that minimize false negatives while simultaneously recovering a large fraction of the periods of objects. In Section 5, we examine both the match fraction to LCDB and this conservative approach that biases against incorrect periods when applied to all objects on the combined data set (SS1∩TESSC1).

4.4. Summary: Comparison of Methods and the Utility of ZTF and TESS Observational Data

In this section, we demonstrated several methods that can be applied to improve the fraction of matches to LCDB, and these methods included σ clipping, using periodogram masking, replacing the derived period in LCDB∩SS1∩TESSC1 with the TESSC1 period when there is a high probability that the TESSC1 period is correct, and using a light-curve amplitude threshold. The fraction of matches to LCDB using the heatmap analysis

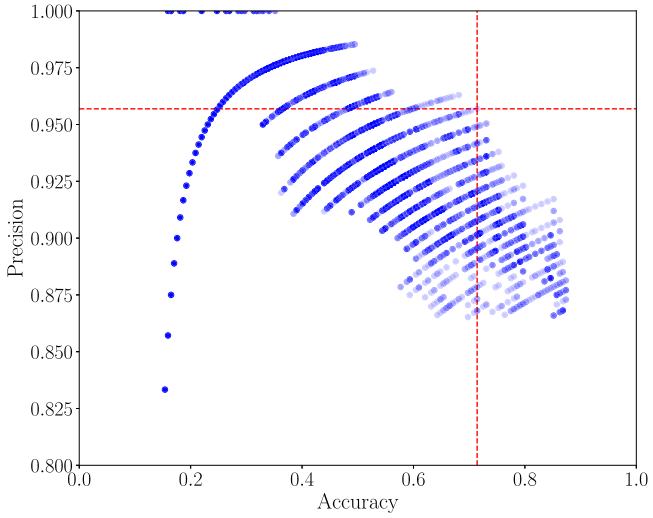

Figure 11. The precision as a function of accuracy for the grid search described in the text. There are a total of 7128 configurations that were searched to find those that optimize the precision metric. Many of the configurations yield the same pair of precision and accuracy values, and so the markers on the plot are translucent such that those with greater intensity denote duplicate configurations. The red dashed lines show an example configuration with a precision of 0.957 and an accuracy of 0.714.

(Figures 9–10) describes when a derived rotation period is likely to be correct; however, it does not bias against assigning an incorrect rotation period to an object. In contrast, the conservative binary classification method biases against assigning an object an incorrect rotation period (a false positive), and so when the classifier reports that an object's classification is incorrect, this should be considered a null result and the period cannot be trusted. The benefit of using the binary classifier over simply using the fraction of matches to LCDB is that there is a higher probability that the classifier will identify correct periods. The drawback of this method is that it favors precision over accuracy, and so several correct rotation periods will be labeled as incorrect (false negatives). In summary, to employ these methods on large catalogs, rather than only reporting a single confidence metric, it is preferable to report the total match fraction, the match fraction as a function of LSP power and number of observations (Figures 9–10), and the conservative binary classification approach that reports correct rotation periods with very high probability.

Recall from Section 4.1 that the match fraction by simply combining SS1∩TESSC1 is 78.4% and 82.9% for exact and aliased matches, respectively. Therefore, the optimizations summarized in Section 4.2.7 are only able to improve the overall match fraction in a small number of remaining cases, as the ∼80% match fraction is remarkable given the drawbacks of the data sets. The baseline approach shows that the match fraction increases substantially when comparing SS1 and SS1∩TESSC1, whereas the match fraction has a smaller increase when comparing TESSC1 and SS1∩TESSC1 (Table 2). Therefore, the benefits of using TESSC1 are clear, but it is not abundantly clear what the benefits are to adding SS1 to TESSC1.

We discuss why incorporating SS1 is important. Recall from Section 4.1 that TESSC1 is not sensitive to $p < 4$ hr periods because of the 30-minute cadence of the full-frame TESS images. Using the LCDB∩SS1∩TESSC1 data set with optimizations "{1, 3, 5, 7} Combined" with $p_{\min} = 3$ hr (182 total objects), there are a total of 31 objects in LCDB with 3–4 hr rotation periods. Of these 31 objects, TESSC1 correctly

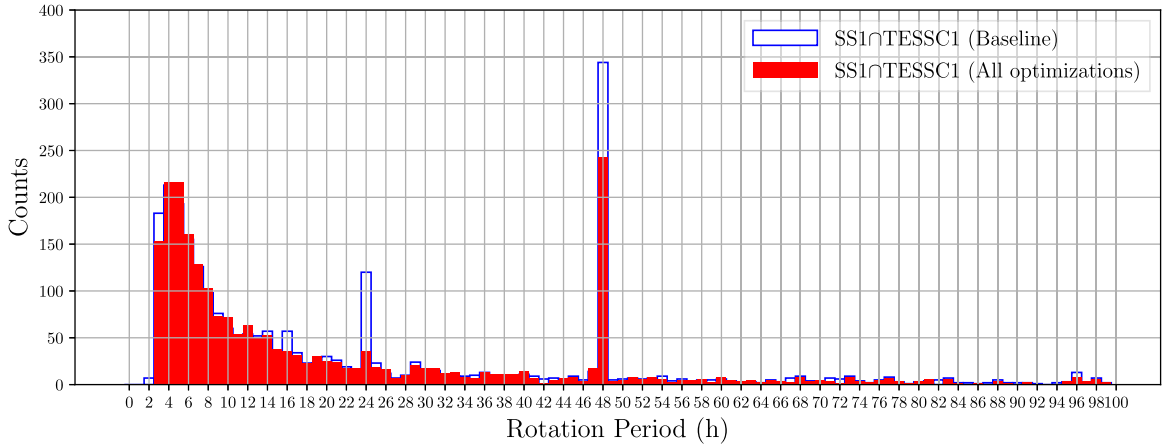

Figure 12. Rotation period distribution of SS1∩TESSC1 comparing the baseline approach to the set of optimizations "{1, 3, 5, 7} Combined" with $p \geqslant 3$ hr from Table 3.

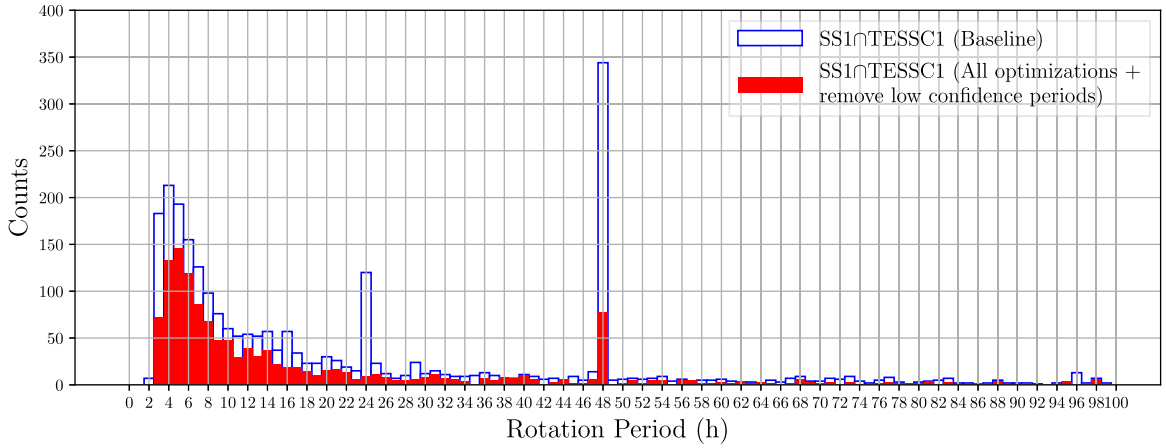

Figure 13. The same as Figure 12, but where low-confidence solutions have been removed.

derives 19 of the rotation periods (considering only exact matches), but there are 6 additional objects that SS1∩TESSC1 derives correctly but TESSC1 derives incorrectly. This shows that incorporating SS1 significantly improves period derivation in this period range.

This trend can be observed in Figure 3, where SS1 recovers more rotation periods within the second bin than TESSC1. It is noteworthy that the true asteroid rotation period distribution has a high fraction of objects with periods in this range, and therefore sensitivity to $p \lesssim 4$ hr is of great importance. This is a concrete example of where SS1 is better than and can complement TESSC1, and in Section 5.4 we show an example light curve for an object that is derived correctly with SS1 but not with TESSC1 with a period of 10.215 hr. To summarize, based on our analysis, SS1 has great utility at $p \sim 2$–4 hr whereas TESSC1 does not, SS1 is sensitive to long period asteroids whereas TESSC1 is not because of its small observing window, and SS1 has moderate utility for intermediate rotation periods compared to TESSC1.

5. Period Derivation on the Combined ZTF and TESS Data Set

Now that we have examined maximizing the rotation period match fraction by comparing our solutions to those reported in LCDB, we apply the combined methods in Section 4.2.7 to the 3168 objects that are included in the SS1∩TESSC1 data set. In particular, we use configuration "{1, 3, 5, 7} Combined" from Table 3 with $p_{\min} = 3$ hr.

Table 4
The Number of Aliases for the Baseline Approach (No Filtering) and Using the Combined Methods "{1, 3, 5, 7} Combined" with $p_{\min} \geqslant 3$ hr in Table 3 and the Combined Method with Low-confidence Period Removal

Alias (hr)	Baseline ($n = 3168$)	Combined ($n = 3008$)	Combined and Low-confidence Removal ($n = 1623$)
16 ± 0.1	57 (0.018)	36 (0.012)	19 (0.012)
24 ± 0.1	120 (0.038)	36 (0.012)	9 (0.006)
48 ± 0.1	344 (0.109)	243 (0.081)	77 (0.047)

Note. The fractions in parentheses refer to the fraction of the sample at the reported alias.

5.1. Rotation Period Distribution

In this section, we examine the following three perspectives on period recovery: (i) using the combined optimizations summarized in Table 3, (ii) using the LSP power in the differential heatmap (Figure 10) to determine which periods are considered reliable, and (iii) using the conservative binary classification approach outlined in Section 4.3.

Table 5
SS1∩TESSC1 Rotation Periods and Amplitudes for $n = 1623$ Objects with High Confidence (∼85%)

MPC Desig.	Rot. Per. (hr)	Amplitude (mag)	Rot. Per. TESSC1	Rot. Per. SS1	TESSC1 Replacement	Class	Rot. Per. LCDB
100052	4.682	0.482	34.946	4.682	N	C	⋯
10026	9.977	0.199	33.726	8.257	N	I	⋯
10041	47.834	0.255	816.327	47.963	N	C	2.564
10065	20.050	0.535	20.047	20.050	N	C	⋯
100683	18.261	0.499	232.480	18.259	N	C	⋯
100727	3.175	0.368	36.920	3.175	N	I	⋯
10080	7.051	0.177	7.052	6.146	N	C	⋯
10117	6.880	0.467	6.881	6.880	N	C	⋯
10128	9.573	0.142	9.571	7.978	N	C	⋯
101632	121.458	1.702	121.458	684.825	Y	C	⋯

Note. We also show the rotation periods for SS1, TESSC1, and LCDB (if applicable) and whether the rotation period was the result of a direct replacement using TESSC1. Column "Class" refers to the result of the conservative binary classifier, where "I" and "C" refer to incorrect and correct class labels, respectively.

(This table is available in its entirety in machine-readable form in the online article.)

Figure 12 plots a histogram of the rotation periods comparing the baseline approach of combining TESSC1 and SS1 (SS1∩TESSC1) to the distribution after applying the optimizations "{1, 3, 5, 7} Combined" with $p \geqslant 3$ hr from Table 3. We observe that we have reduced solutions at the aliases: the 16 hr alias is no longer visible, the 24 hr alias is visible but much reduced, and the 48 hr alias is still pronounced. Table 4 summarizes the number of solutions at aliases in Figure 12.

Recall that there are only $n_{\mathrm{LCDB}} = 206$ objects in LCDB that overlap with our sample (Table 3). Thus, if the $n = 3168$ objects in SS1∩TESSC1 were drawn from the same distribution as the $n_{\mathrm{LCDB}} = 206$ objects in LCDB∩SS1∩TESSC1, then the match fraction values in the cumulative heatmap (Figure 9) would be directly applicable to the objects in SS1∩TESSC1. However, we know from the histogram shown in Figure 12 that our period distribution is not the same, which is clearly indicated by the overabundance of aliased solutions. Thus, the differential heatmap (Figure 10) is a better statistical tool for determining whether we should be confident in a given solution despite the abovementioned challenges regarding small number statistics.

Consider from Figure 10(b) that objects with an LSP power <0.333 are unreliable and objects with an LSP power $\geqslant 0.333$ have a minimum confidence level of 0.848. Figure 13 shows the same histogram as Figure 12, where all periods with an LSP power <0.333 are removed from the sample. This yields a total of $n = 1623$ objects, indicating that 1545 period solutions were removed because they have a low LSP power. For these $n = 1623$ objects, the minimum confidence that we assign is 0.848. Comparing Figure 12 to Figure 13 shows that many of the 1545 (removed) objects with an LSP power <0.333 are those that were located in the 48 hr bin (11% of the 1545 removed objects). We find that combining the optimizations and removing low-confidence solutions is unable to completely eliminate aliased period solutions at 48 hr, but it is much improved over both the baseline and the optimization set "{1, 3, 5, 7} Combined" that did not remove low-confidence period solutions from the sample.

Section 4.3 used a binary classification approach to label the rotation periods as either incorrect or correct based on amplitude, number of observations, and LSP power. We find that there are a total of 897 asteroids assigned a correct label and 2271 assigned an incorrect label. The asteroids assigned a correct label are plotted in Figure 14. We find that, compared to Figures 12 and 13, there is only a slight overabundance of periods at 48 hr, indicating that the binary classifier filters out many of the aliased solutions. Recall that this conservative approach trades precision for accuracy by penalizing false positives, and so the total number of objects recovered is much lower than in Figure 13, as these objects have a higher probability of being assigned correct rotation periods.

Table 5 reports the rotation periods of the $n = 1623$ asteroids defined where all periods with an LSP power <0.333 were removed (Figure 13). For each of these objects we also report the binary classification prediction, where 897 were labeled as correct and 726 were labeled as incorrect.

5.2. Rotation Period Comparison with TESSC1 and SS1

We examine the rotation period distributions of SS1 as a function of the SS1∩TESSC1 rotation period in the left panel of Figure 15. While a substantial fraction of periods are in agreement, we clearly observe the aliased solutions at 48 hr in SS1 that are not in agreement with SS1∩TESSC1. Examining the SS1∩TESSC1 amplitudes of these objects, we find that they are generally fairly low at 48 hr, as most are <0.25 mag. We also observe the effect of pseudo-aliases (J. T. VanderPlas 2018), which appear as the curved triangle shape in the left panel. These are due to interactions between an object's real period and the strongest alias in the periodogram (J. Ďurech et al. 2022; D. E. Trilling et al. 2023b).

The right panel of Figure 15 plots the rotation periods of TESSC1 as a function of SS1∩TESSC1. Here the maximum rotation period for TESSC1 is ∼816 hr, which is why none of the periods in TESSC1 exceed this value. We find that there are an overabundance of aliases clearly observed at 48 hr in SS1∩TESSC1, which is unsurprising since we showed this in Figure 13 and in Table 4. However, of great interest is the observation that the pseudo-aliases are no longer visible when comparing TESSC1 to SS1∩TESSC1. This is because TESSC1 does not have aliases, and so the pseudo-aliases are eliminated.

5.3. Light-curve Amplitude Distribution

Figure 16 plots the amplitudes of the $n = 1623$ sample of objects in SS1∩TESSC1. The figure is consistent with Figure 6(a), except that there are a smaller fraction of objects

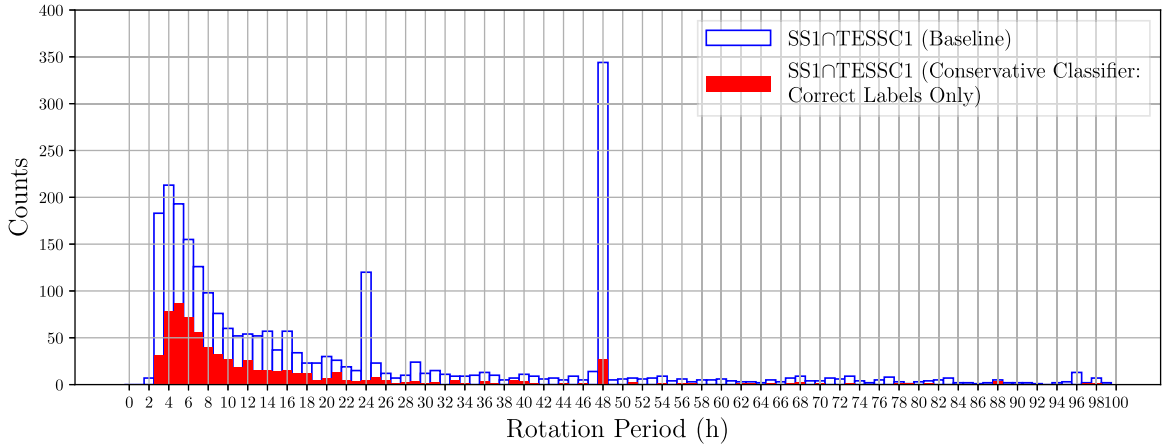

Figure 14. The same as Figure 12, but where we only show those classified as correct with the conservative binary classifier described in Section 4.3. The thresholds are as follows: an amplitude of 0.125 mag, 120 observations, and an LSP power of 0.5.

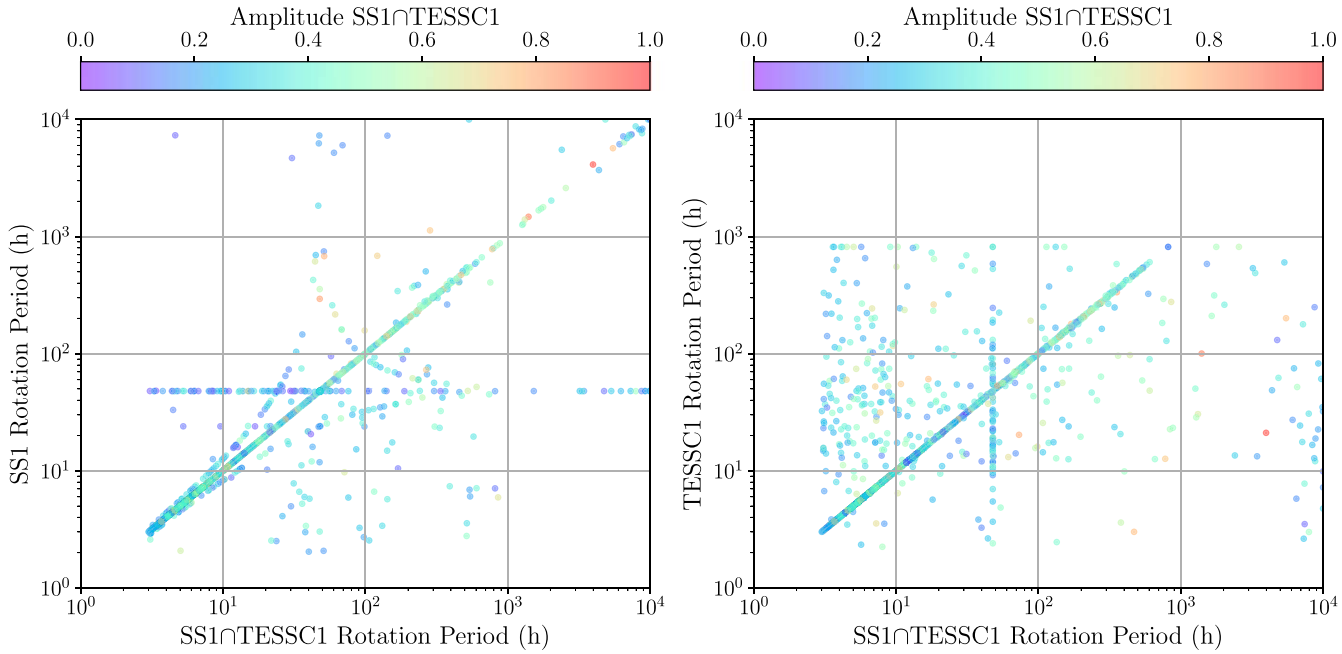

Figure 15. Comparison of rotation period distributions as scatterplots for SS1 and TESSC1 as a function of SS1∩TESSC1, plotted in the left and right panels, respectively. The SS1∩TESSC1 rotation periods and color-coded amplitudes are from the $n = 1623$ sample that has all optimizations applied and low-confidence solutions removed. In contrast to prior figures, such as Figure 12, the plot is shown on a log scale such that the entire period range can be shown. Consistent periods between the two data sets appear on the diagonal line.

with low amplitudes, as we discarded many of them to maximize the fraction of matches in the sample. The shape of an asteroid can be constrained through $\Delta\text{mag} = -2.5 \log \frac{b}{a}$, where b and a are the minor- and major-axis lengths of the ellipsoidal asteroid shape, respectively. There are few objects with $\frac{b}{a} \geqslant 0.9$, and a single asteroid cannot be this elongated; therefore, these asteroids are likely to be binary systems.

5.4. Example Light Curves and Implications

In all prior sections we described facets of light-curve generation and fidelity without showing any example light curves. In this section we show a selection of interesting light curves from the $n = 1623$ sample of objects in SS1∩TESSC1. In the plots, we show the derived light curve of SS1∩TESSC1 without TESSC1 replacement (if applicable), such that we can observe the light curves generated by the combined data set in each instance.

Object 339: Figure 17 shows the light curves for object 339 for SS1∩TESSC1 ($p = 2349.795$ hr), TESSC1 ($p = 5.967$ hr), and SS1 ($p = 47.957$ hr) and the unphased SS1∩TESSC1 data with error bars ordered from top to bottom, respectively. The light curve for SS1∩TESSC1 (top panel) has a rotation period of $p = 2349.795$ hr, but TESSC1 is not sensitive to rotation periods $\gtrsim$816 hr, and so the light curve is poorly sampled after $\sim$600 hr (i.e., t_{window} for this object). The derived rotation period is a poor solution for this object. The second panel shows the data for TESSC1, and we observe an excellent fit to the data. Here the rotation period is $p = 5.967$ hr and LCDB reports $p = 5.974$ hr, so the rotation periods are in agreement. The third panel shows the light curve generated with SS1, and we observe a very poor fit to the data. This light curve has a

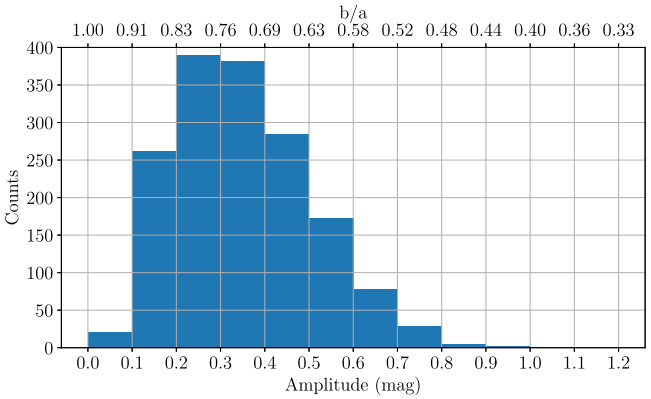

Figure 16. The amplitudes of the $n = 1623$ sample of objects in SS1∩TESSC1 that have all optimizations applied and low-confidence solutions removed. The top horizontal axis shows the $\frac{b}{a}$ ratio, which is calculated directly from the amplitude (Δmag) where $\Delta\text{mag} = -2.5 \log \frac{b}{a}$.

rotation period of 47.957 hr and is an example of a quintessential alias where the only signal that can be detected by LSP is a function of the diurnal observing schedule. Using the optimizations described in the prior sections, the SS1∩TESSC1 period of $p = 2349.795$ hr was replaced by the TESSC1 $p = 5.967$ hr, which is the correct period. The bottom panel shows the unphased SS1∩TESSC1 data, where we observe that some of the SS1 observations temporally overlap the TESSC1 data, but most of the SS1 data are in a separate apparition.

Object 18975: Figure 18 reports the light curves for object 18975. This object is not reported in LCDB. The rotation periods are $p = 404.538$ hr, $p = 413.212$ hr, and $p = 404.608$ hr for SS1∩TESSC1, TESSC1, and SS1, respectively. All of the rotation periods are in agreement, and the object appears to be a slowly rotating asteroid. This set of light curves is interesting because if we consider the TESSC1 (middle panel) and SS1 (bottom panel) light curves on their own, one may be skeptical that the derived rotation periods are accurate because the light curves are poorly sampled (i.e., there are large gaps in the folded light curve where no observations are present). However, the light curve is more robustly sampled in SS1∩TESSC1 (top panel), and so this is an example where combining the two data sets allows us to have greater confidence in the solution. In our sample of $n = 1623$ objects there are numerous examples of this, including the case where the three period solutions disagree owing to poor sampling of the folded TESSC1 and SS1 light curves, but the SS1∩TESSC1 light curve is more compelling owing to better sampling.

Object 363: Figure 19 is interesting for two reasons. First, SS1 (bottom panel) generates a light curve that is clearly very poorly sampled; however, we find that TESSC1 (middle panel) has an excellent fit of $p = 16.801$ hr. When we examine SS1∩TESSC1 (top panel), we find that for this rotation period of $p = 16.811$ hr the SS1 observations clearly do not contribute to the goodness of the fit. In other words, using only the TESSC1 observational record produces a better light curve than SS1∩TESSC1, but the solution provided by SS1∩TESSC1 yields a good fit despite the noise imparted by the SS1 observations. The second interesting observation is that the rotation period for TESSC1 (middle panel) is $p = 16.801$ hr and the reported rotation period in LCDB is $p = 8.401$ hr, which is a factor 0.5 of our solution. This is considered an aliased match, where the derived rotation period is a factor of two greater than the rotation period reported in the literature. We attribute this discrepancy to the very low amplitude of the object, which is 0.083 mag using the TESSC1 rotation period, and as shown in Figure 8, low amplitudes may lead to rotation period matches that are incorrect by a factor 0.5 or 2 of the rotation period reported by LCDB.

Object 14376: Figure 20 shows that all three light curves are in agreement, as they all report a rotation period of $p \approx 5.85$ hr. Based on a manual inspection of the light curves, it appears that the rotation period of $p \approx 5.85$ hr is plausible. However, LCDB reports $p = 5.614$ hr, which is in disagreement with our solution since we consider that rotation periods are matching if they are within 3% of each other. Thus, while this fit is clearly reasonable, it is deemed incorrect in our analysis. This example shows that while many derived light-curve solutions are incorrect for obvious reasons, such as poor sampling, even light curves with good fits may be incorrect. It is possible that the solution in LCDB is incorrect, or another interesting possibility is that the rotation period has changed. This further demonstrates the challenges of reliably deriving rotation periods from sparse or dense photometry.

Object 3176: In the prior light-curve figures we clearly observe that on average there are more data points in the observational records of TESSC1 than in SS1. Additional observations for an object help constrain the rotation period because it reduces sampling sparsity. Figure 21 shows an example where SS1∩TESSC1 and TESSC1 are in agreement and yield rotation periods of $p = 20.457$ hr and $p = 20.433$ hr, respectively. However, the SS1 rotation period of $p = 10.224$ hr is an exact match to that reported in LCDB, which is $p = 10.215$ hr (the SS1∩TESSC1 and TESSC1 rotation periods are aliased matches, as they are a factor 2 greater than the LCDB rotation period). This demonstrates that while fitting light curves with additional observations is generally beneficial, there are cases where fewer observations may yield a better derived rotation period. We also note that while the $p = 10.224$ hr rotation period for SS1 (bottom panel) is clearly a good fit, the $p = 20.457$ hr rotation period (top panel) is not a good fit to the SS1 data.

5.5. Discussion

We examined several methods for improving derived period fidelity when combining observational records from ground- and space-based facilities. We find that combining the data records from ZTF and TESS improves the overall match fraction with LCDB, and we applied the optimizations on the sample of objects in LCDB∩SS1∩TESSC1 to the much larger SS1∩TESSC1 sample. In this section, we discuss implications for forthcoming Rubin Observatory operations, the use of the binary classification approach on other catalogs, and other optimizations beyond the scope of this paper.

5.5.1. Rubin Observatory

In the context of preparing for the Rubin Observatory, several of the optimizations and insights may need to be reexamined. We outline these as follows:

1. The amplitude optimization should be reexamined, as Rubin will be more sensitive to amplitudes <0.075 mag.

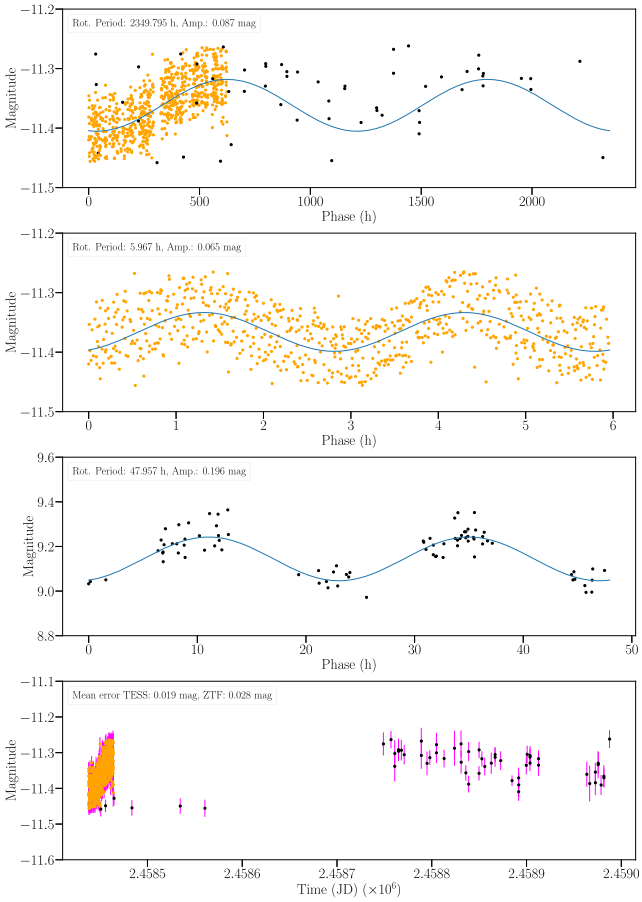

Figure 17. Light curves of object 339. Panels from top to bottom are as follows: SS1∩TESSC1 ($p = 2349.795$ hr), TESSC1 ($p = 5.967$ hr), SS1 ($p = 47.957$ hr), and the unphased data for SS1∩TESSC1. Orange and black markers refer to observations from TESSC1 and SS1, respectively. The bottom panel shows photometric errors that are plotted using magenta bars, and the markers are plotted on top of the error bars to prevent them from being obfuscated.

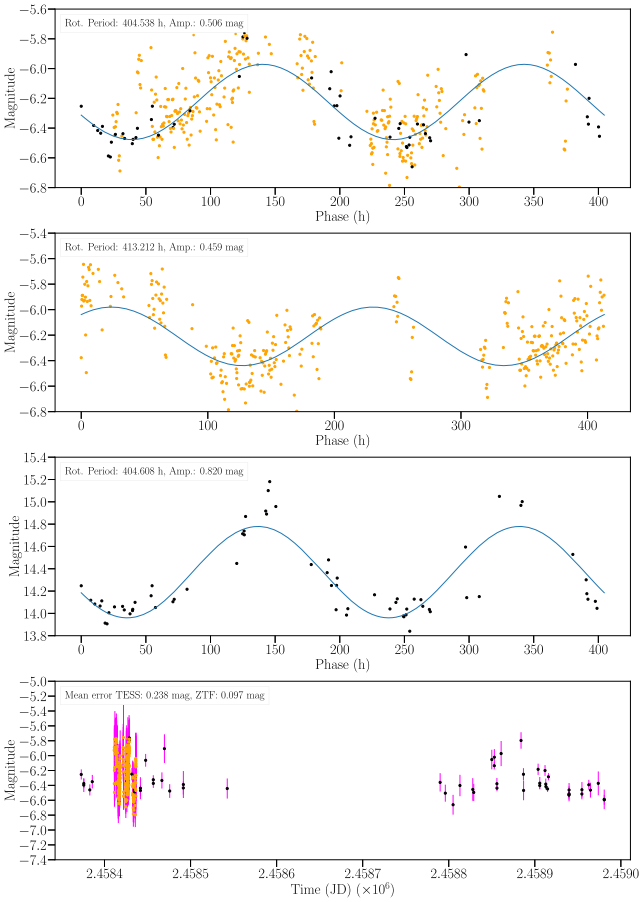

Figure 18. Light curves of object 18975. Panels from top to bottom are as follows: SS1∩TESSC1 ($p = 404.538$ hr), TESSC1 ($p = 413.212$ hr), SS1 ($p = 404.608$ hr), and the unphased data for SS1∩TESSC1.

2. There are a number of light curves that we found to be nonsinusoidal (see, e.g., Figure 19). These light curves may be better fit by other period-finding algorithms that are not constrained to sinusoids.
3. We examined period replacement with TESSC1 and SS1 and found that the former was useful but the latter was not. We found that few objects in SS1 had a sufficient number of observations and LSP power (Figure 5) to achieve an 85% match fraction to LCDB. The next SNAPS data release will provide many more ZTF observations per object, and so period replacement with this new data set should be reexamined if augmenting Rubin data with ZTF data.
4. The binary classification approach will need to be reexamined to select a configuration that optimizes for precision and accuracy. For example, Rubin will have lower photometric error, and so the observation threshold may be reduced compared to that used in this paper.

Furthermore, there are two implications for the Rubin Observatory that originate from the analysis of Figure 21 that we outline as follows.

1. While having a larger number of observations for an object reduces the sparsity of the folded light curve, which improves the fit, another factor that may be of great importance is having better-quality (lower photometric uncertainty) observations. D. Kramer et al. (2023a) showed that for the Rubin Observatory an observing cadence that favors longer exposures and thus reduces the total number of observations for each object yields a higher rotation period match fraction to LCDB compared to a cadence with shorter exposures and more observations per object.
2. Because we find that a moderate number of observations (Figure 21 showing SS1) can yield accurate rotation periods, the Rubin Observatory will produce high-fidelity light curves within the first few years of operation. Thus, a major task will be to automate the characterization of light-curve fidelity to curate light curves that have a high probability of having correct rotation periods and biasing against reporting incorrect rotation periods.

In the context of examining ZTF and SSPDB data, D. Kramer et al. (2023b) examined several methods for improving the overall rotation period match rate to LCDB, and they discuss the prospect of combining optimizations. We find that combining optimizations is only beneficial for capturing rotation periods in a small number of instances and instead combining the two data sets to produce SS1∩TESSC1 had a larger overall benefit on period fidelity. One caveat to this analysis is that the LCDB is not a particularly large catalog, and so we were only able to compare against 222 objects.

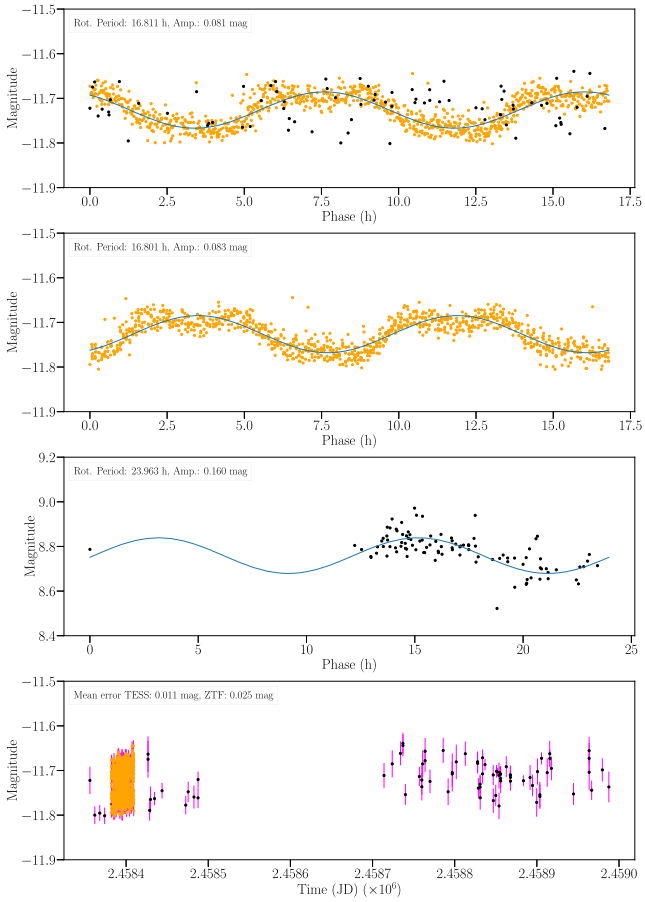

Figure 19. Light curves of object 363. Panels from top to bottom are as follows: SS1∩TESSC1 ($p = 16.811$ hr), TESSC1 ($p = 16.801$ hr), SS1 ($p = 23.963$ hr), and the unphased data for SS1∩TESSC1. The LCDB rotation period is reported as 8.401 hr, which is a factor 0.5 of the TESSC1 derived rotation period solution of 16.801 hr (middle panel). This light curve is also remarkable because it has extra structure at the peaks instead of simply having a sinusoidal shape.

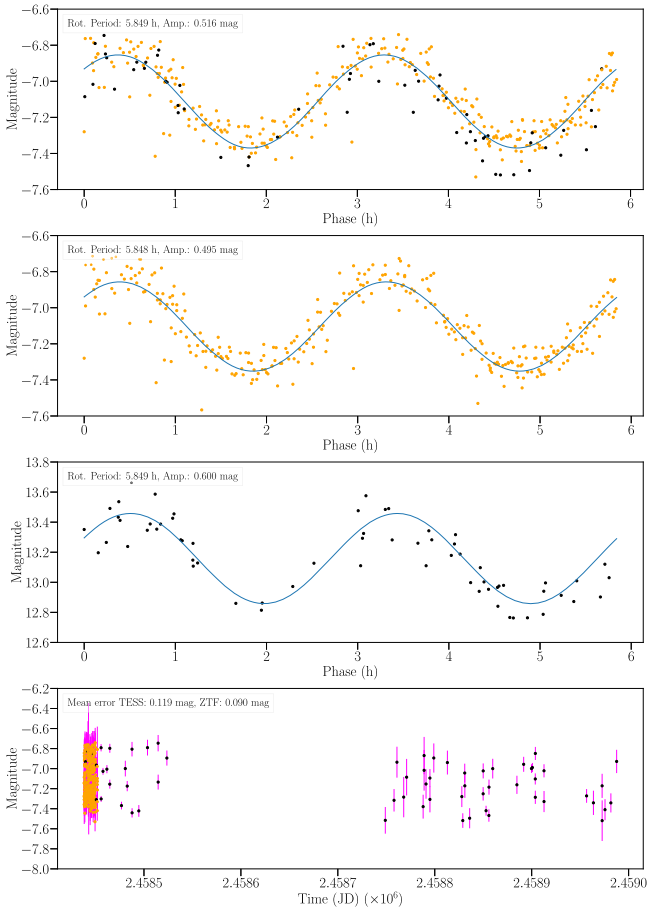

Figure 20. Light curves of object 14376. Panels from top to bottom are as follows: SS1∩TESSC1 ($p = 5.849$ hr), TESSC1 ($p = 5.848$ hr), SS1 ($p = 5.849$ hr), and the unphased data for SS1∩TESSC1.

Therefore, a larger reference sample would help us better understand the utility of the optimizations presented here.

It is clear that low-amplitude asteroids are often those where LSP predicts a period at an alias, and for ZTF these appear at $p \in \{16, 24, 48\}$ hr. Thus, these asteroids fundamentally cannot produce a light curve and are simply intractable for period derivation, which implies that the ≈85% match fraction to LCDB achieved in Table 3 may be close to the upper limit on derived period fidelity assuming reasonable constraints on the minimum amplitude and minimum rotation period.

5.5.2. Binary Classification Approach: Application to Other Catalogs

The binary classification approach employed here is used to bias against reporting incorrect asteroid rotation period solutions. The method can be applied to other catalogs where there is a baseline (or fiducial) set of correct rotation period solutions and another set of solutions of unknown quality. In this paper, our baseline data set was LCDB.

While we selected binary classification labels using thresholds for the LSP power, light-curve amplitude, and number of observations, these parameters will need to be reevaluated based on the input data set of unknown quality. For example, observations obtained with lower photometric error may require lower observation and amplitude thresholds to obtain similar accuracy and precision levels reported in Section 4.3. Furthermore, it may be preferable to select accuracy and precision based on the receiver operator characteristic curves, which illustrates the relationship between the true-positive and false-positive rates.

Deriving a generalizable set of parameter thresholds for the binary classification approach is not possible, as we only considered two surveys in this paper and so the parameters were tailored to those surveys. However, a promising future research direction is to combine data from multiple surveys. If the surveys span a large range of parameter values, including mean photometric error and number of observations per object, it may be possible to extract generalizable parameters that yield high precision with minimal loss in accuracy. This research direction may benefit from a supervised machine learning approach, as the number of parameters needed will exceed the three parameters investigated here.

5.5.3. Other Optimizations

There may be several other avenues of investigation that may improve period fidelity. We focused on combining several well-known methods in Section 4 but also attempted two other less obvious methods that were not useful, and so we did not report these results.

We know that LSP is good for fitting the light curves of asteroids because they are generally sinusoidal in shape. We

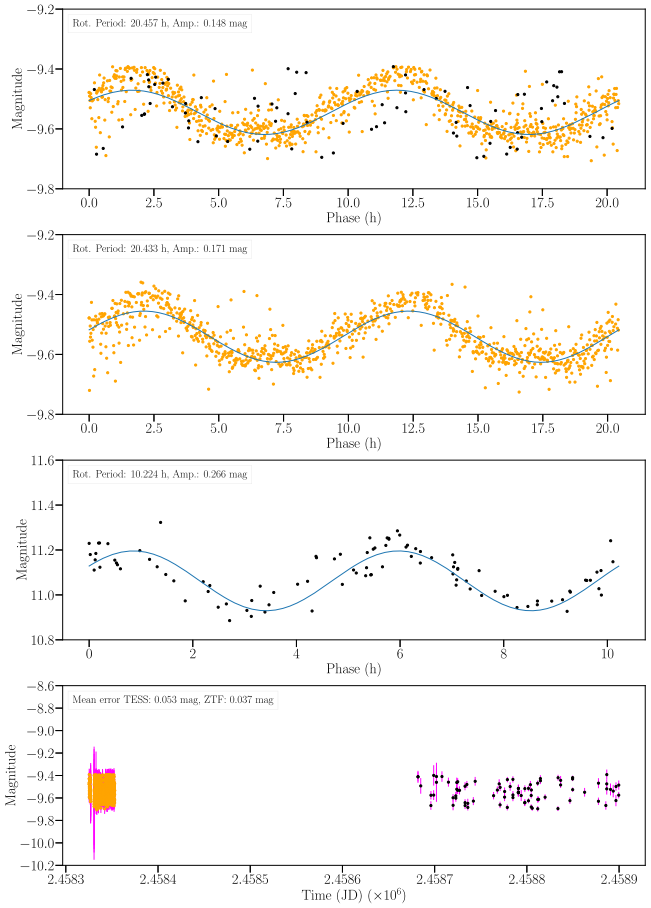

Figure 21. Light curves of object 3176. Panels from top to bottom are as follows: SS1∩TESSC1 ($p = 20.457$ hr), TESSC1 ($p = 20.433$ hr), SS1 ($p = 10.224$ hr), and the unphased data for SS1∩TESSC1.

examined using the SUPERSMOOTHER (J. H. Friedman 1984) algorithm, which takes advantage of structure in the data that can be used to disambiguate periods detected near aliases. For example, M. Gowanlock et al. (2022) found that SUPERSMOOTHER was better than LSP for computing the periods of RR Lyrae stars. However, we found that LSP is better than SUPERSMOOTHER for computing the rotation periods of asteroids. Consequently, this method was not presented here.

Another approach to disambiguate aliases is to use Monte Carlo sampling of the observational record to produce a set of periodograms for each object. Combining the set of periodograms to select the frequency with the greatest combined power may reduce aliases (D. Kramer et al. 2023b). We selected a sample of objects for which our methods were unable to derive the correct rotation period where most of these objects had aliased period solutions. However, we found that the Monte Carlo approach was unable to improve period fidelity in these instances. It is likely that this is because of the low-amplitude problem that we outlined above—many of these objects have very low amplitudes for which producing a light curve is intractable.

6. Prospects for the Rubin Observatory

In the prior sections we examined several approaches to improve the fidelity of derived asteroid rotation periods using several optimizations in the context of combining data from ground- and space-based observatories. In this section, we

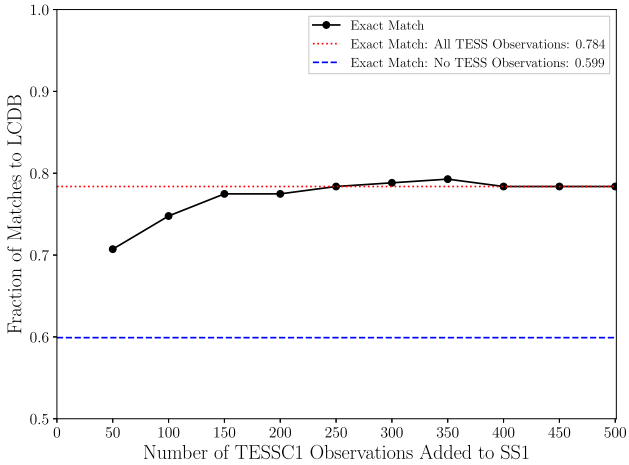

Figure 22. The exact match fraction to LCDB as a function of the number of TESSC1 observations added to SS1. The top horizontal line refers to the exact match fraction when all TESSC1 observations are added to SS1, and the bottom horizontal line refers to when there are no TESSC1 observations added to SS1. These exact match fractions are 0.784 and 0.599, respectively (Table 2).

examine a practical aspect of Rubin Observatory operations that is driven by the following question: to what extent can a companion (ground- or space-based) observatory significantly improve light-curve fidelity and derived rotation periods of regularly varying sources observed with the Rubin Observatory, and how much telescope time is required?

Despite the impressive capabilities of the Rubin Observatory, it will still be limited by the diurnal observing schedule, which will produce aliases. By employing a companion telescope to augment the LSST catalog, aliases may be partially suppressed (as we observed examining the SS1∩TESSC1 data set). Because telescopic observations are expensive, it would be interesting to know how much observing time on a companion telescope is necessary to eliminate a substantial fraction of the rotation periods reported at the aliases generated by the Rubin Observatory's schedule and improve the overall yield of correct rotation period solutions.

6.1. A Future TESS-like Observatory

Here we explore the benefit of including data from a future TESS-like observatory (which could indeed be TESS in an extended mission). To assess the above question, we use the $n = 222$ objects in LCDB∩SS1∩TESSC1. We use the entire observational record for each object in SS1 (thus using ZTF data to simulate the LSST catalog) and then incrementally add TESSC1 to SS1, where we are simulating using TESS (or a TESS-like facility) as a companion telescope to the Rubin Observatory. We do not apply any of the optimizations outlined in Section 4, as we are simulating incrementally adding observations to the ground-based observational records. As such, some of the optimizations would not be possible, such as knowing which observations to σ-clip without the full observational record for an object (e.g., with a small observational record for an object, we may inadvertently remove observations that are useful for fitting the light curve).

The number of observations for each object in TESSC1 varies. Therefore, when we combine the data between SS1 and TESSC1, we use either the desired number of observations that we wish to add to SS1 or the total number of observations for

an object, whichever is smaller. For example, if we want to use 100 TESSC1 observations for an object but the observational record only contains 90 entries for that object, then we are limited by the 90 observations and use the full observational record containing these observations.

Figure 22 plots the exact match fraction to LCDB as a function of the number of observations added per object from TESSC1 to SS1. The upper and lower bound match fractions are given by the two horizontal lines, which refer to when all (or none) of the TESSC1 observations are added to SS1. These fractions are reported in Table 2, where we find match fractions of 0.784 and 0.599, respectively.

We find that after adding roughly 150 TESSC1 observations to SS1 we obtain a match fraction approaching the upper bound of 0.784. Given that TESS uses a 30-minute full-frame image cadence, 150 observations/full frames would require roughly 3 days of TESS observations.

As a targeted companion facility that is synchronized to LSST observations, this is an impractical solution, as LSST covers most of the available southern sky during a 3-day period, whereas TESS only covers a single field of view—around 10% of one hemisphere—during this same time period. However, over the course of a year, TESS covers an entire hemisphere with >3 days of continuous coverage, and LSST also covers an entire hemisphere of the sky. The TESS-like data need not be contemporaneous with LSST-like observations; all that matters is the uninterrupted cadence that complements the LSST-like observations. Therefore, if TESS were observing during LSST's survey, the alias solutions would be much reduced and the overall match fraction increased as a natural consequence of combining these data sets with two different cadences. (If this future TESS uses the 30-minute cadence, then we still have no new power to address potential periods less than 3 hr.)

The above estimates may have underpredicted the delivered performance, though. Consider that the SS1 data set derived from ZTF is roughly 10% of the scale of LSST (D. E. Trilling et al. 2023b). Thus, if the same experiment were to be carried out with the Rubin Observatory, we would expect that the match fraction without any companion telescope data would be higher than the 78.4% reported here because each object would be observed more with the Rubin Observatory than ZTF over the same time period. In addition, recall from Section 4.2.6 that it is challenging to derive the rotation periods of asteroids with low light-curve amplitudes. The Rubin Observatory will have lower photometric error ($\leqslant$0.1 mag; S. Navarro-Meza et al. 2021) than ZTF, and so we expect that it will yield many more low-amplitude objects with light-curve amplitudes <0.1 mag that have correct rotation period solutions. Given that we expect many of the objects for which we cannot derive a correct rotation period solution to be those with low amplitudes, we believe that this capability will dramatically improve derived rotation period fidelity.

We therefore conclude that if TESS is still operational in 2025 and after, combining TESS data with LSST data will be extremely useful for increasing the fidelity of period solutions.

ESA's Euclid mission (R. Laureijs et al. 2011; G. D. Racca et al. 2016) complements both TESS and LSST observations, and the joint Rubin-Euclid working group has already offered some recommendations about how the combined data sets may enhance the science return of both projects (L. P. Guy et al. 2022). The footprint of Euclid's Wide Survey avoids the ecliptic, but there are plenty of asteroids that will still be detected in Euclid data. The imaging sequence of Euclid will provide several images over 1 hr (Euclid Collaboration 2022). Thus, Euclid data may be useful to identify asteroids with short rotation periods, which neither TESS nor LSST will be very successful at identifying. However, in general Euclid will not offer enough complementary data to significantly improve a large number of LSST period solutions.

6.2. Augmenting LSST Data with Ground-based Data

Additional data from the Rubin Observatory itself might be useful to resolve degeneracies through data from the Deep Drilling Fields (DDF). Not all DDF observations have been defined at this point, but it is likely that at least some of these DDF campaigns will be a continuous all-night stare at a single LSST pointing. This is a TESS-like (i.e., continuous) set of perhaps ~1000 observations obtained over around 10 hr. This may help resolve degeneracies at the shortest periods, but it is not likely to address diurnal aliases. In a future investigation we will model the performance of LSST period solutions in the presence of both the standard "wide-fast-deep" cadence (which is ZTF-like) and DDF observations.

It is possible that other ground-based observatories or networks that have broad longitudinal coverage, such as the ATLAS network (J. L. Tonry et al. 2018), which currently has four operational sites (one site in South Africa, one in Chile, and two in Hawai'i), could collectively observe the same objects without a break in day−night observing. This would help reduce the number of rotation period solutions at aliases and would be less expensive than pursing a companion space-based telescope for LSST. However, ATLAS's limiting magnitude of perhaps 20 is not well matched to LSST's limiting magnitude of 24.5; the vast majority of sources observed by LSST will be too faint for ATLAS.

In summary, there are no obvious existing or near-term ground-based assets that will provide the data necessary to improve period solutions on a large scale, but both LSST and ATLAS (and perhaps others) may provide limited capacity to improve solutions derived only from LSST wide-fast-deep data.

7. Conclusion

This paper examined improving the derived asteroid rotation period fidelity when combining data from ground- and space-based facilities. In particular, we used the SS1 (D. E. Trilling et al. 2023b) and TESSC1 (A. McNeill et al. 2023) data sets, which were collected with ZTF and TESS, respectively. We found that our baseline approach of combining the two data sets had the greatest benefit on improving our rotation period match fraction to LCDB. Combining several optimizations further improved the match fraction in several instances; however, the impact of the optimizations has diminishing returns compared to simply combining the two data sets together. This is because many of the objects that had incorrect rotation periods were those that are well-known to be challenging to derive, such as those with low light-curve amplitudes that often have rotation periods that are reported at aliases. Consequently, we find that the ≈85% match fraction may be close to the upper limit for this data set because many of the remaining ≈15% of asteroids may have a spherical

morphology or pole-on orientation that makes light-curve derivation intractable.

The SS1 data set is limited by 0.1 mag photometric error, and so we are not sensitive to amplitudes less than this value. The TESSC1 data set has even greater photometric error than SS1, so it is even less sensitive to low-amplitude (spherical) asteroids. The low-amplitude problem described above will improve when the Rubin Observatory is online, as it will have lower photometric error than ZTF or TESS.

We also discussed Rubin Observatory operations and the prospects of augmenting the LSST data stream with data from space- or ground-based observatories to reduce the fraction of aliases. While none of the solutions for eliminating aliases are straightforward (or necessarily likely such as using TESS in 2025 and thereafter), the analysis points to several possible avenues that can be pursued over the 10 yr life span of Rubin Observatory operations.

Recall that the $\gtrsim$0.1 mag photometric error yielded by ZTF and TESS yields inaccurate period solutions when light-curve amplitudes are $<$0.1 mag. An interesting future work direction that we will address is testing the "sphere-abundant" and "sphere-limited" hypotheses illustrated in Figure 7 (R. Strauss et al. 2024, in preparation). This future work direction will be conducted as part of the DECam Ecliptic Exploration Project (DEEP; D. E. Trilling et al. 2023a), which will yield observations with a mean photometric error of roughly 0.03 mag. This is significantly lower than the $\sim$0.1 mag photometric error provided by ZTF and will allow us to determine which of the two hypotheses above are plausible.

Other future work includes incorporating data from other sources, such as observations of asteroids from Kepler (A. Sergeyev et al. 2023) to create a unified observational record for analysis by our team and the community. Another direction is to examine accuracy as a function of precision for the binary classification approach to deriving rotation periods when using data from numerous sources. This direction may yield insights into the fundamental parameters needed to correctly derive (with very high probability) the rotation periods of asteroids using heterogeneous data sources. We will also incorporate these analyses into the Solar System Notification Alert Processing System (SNAPS) alert broker infrastructure to further enhance our derived data products for the LSST era.

Acknowledgments

We thank the anonymous referee for a very thorough review of the manuscript which significantly strengthened the paper. This material is based on work supported by the National Science Foundation under grant Nos. 2042155 and 2206796. We acknowledge support by the Arizona Board of Regents' Regents Innovation Fund and the Technology Research Initiative Research Fund (TRIF) Small Research Equipment Acquisition Program (SREAP) at Northern Arizona University. We thank Siegfried Eggl for fruitful discussions regarding the combination of data from space- and ground-based observatories. We thank Ryder Strauss for providing the estimated photometric error for the DEEP project.

ZTF is a public-private partnership, with equal support from the ZTF Partnership and from the US National Science Foundation through the Mid-Scale Innovations Program (MSIP). The ZTF partnership is a consortium of the following universities and institutions (listed in descending longitude): TANGO Consortium of Taiwan; Weizmann Institute of Sciences, Israel; Oskar Klein Center, Stockholm University, Sweden; Deutsches Elektronen-Synchrotron & Humboldt University, Germany; Ruhr University, Germany; Institut national de physique nucléaire et de physique des particules, France; University of Warwick, UK; Trinity College, Dublin, Ireland; University of Maryland, College Park, USA; Northwestern University, Evanston, USA; University of Wisconsin, Milwaukee, USA; Lawrence Livermore National Laboratory, USA; IPAC, Caltech, USA; Caltech, USA.

Software: Astropy (T. P. Robitaille et al. 2013), SciPy (P. Virtanen et al. 2020), NumPy (C. R. Harris et al. 2020), and pandas (pandas development team 2020).

Facility: TESS, ZTF

ORCID iDs

Michael Gowanlock https://orcid.org/0000-0002-0826-6204
David E. Trilling https://orcid.org/0000-0003-4580-3790
Andrew McNeill https://orcid.org/0009-0005-9955-1500
Daniel Kramer https://orcid.org/0000-0002-6676-1713
Maria Chernyavskaya https://orcid.org/0000-0002-6292-9056

References

Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, PASP, 131, 018002
Bernardinelli, P. H., Bernstein, G. M., Jindal, N., et al. 2023, ApJS, 269, 18
Cabrera-Vives, G., Reyes, I., Förster, F., Estévez, P. A., & Maureira, J.-C. 2017, ApJ, 836, 97
Coughlin, M. W., Burdge, K., Duev, D. A., et al. 2021, MNRAS, 505, 2954
Drake, A. J., Catelan, M., Djorgovski, S. G., et al. 2013, ApJ, 763, 32
Duev, D. A., Mahabal, A., Masci, F. J., et al. 2019, MNRAS, 489, 3582
Ďurech, J., Vávra, M., Vančo, R., & Erasmus, N. 2022, FrASS, 9, 809771
Erasmus, N., Kramer, D., McNeill, A., et al. 2021, MNRAS, 506, 3872
Euclid Collaboration, Scaramella, R., Amiaux, J., et al. 2022, A&A, 662, A112
Friedman, J. H. 1984, Technical Report SLAC-PUB-3477, STAN-LCS 005, SLAC National Accelerator Laboratory (SLAC), Menlo Park, CA (United States)
Gowanlock, M., Butler, N. R., Trilling, D. E., & McNeill, A. 2022, A&C, 38, 100511
Gowanlock, M., Kramer, D., Trilling, D. E., Butler, N. R., & Donnelly, B. 2021, A&C, 36, 100472
Graham, M. J., Drake, A. J., Djorgovski, S. G., et al. 2013, MNRAS, 434, 3423
Guy, L. P., Cuillandre, J.-C., Bachelet, E., et al. 2022, Rubin-Euclid Derived Data Products: Initial Recommendations, v.2.0.0, Zenodo, doi:10.5281/zenodo.5836022
Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Natur, 585, 357
Juric, M., Eggl, S., Jones, L., et al. 2021, BAAS, 53, 101.06, https://baas.aas.org/pub/2021n7i101p06
Kramer, D., Gowanlock, M., & Trilling, D. 2023a, AJ, submitted
Kramer, D., Gowanlock, M., Trilling, D., McNeill, A., & Erasmus, N. 2023b, A&C, 44, 100711
Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, arXiv:1110.3193
Lomb, N. R. 1976, Ap&SS, 39, 447
Matheson, T., Stubens, C., Wolf, N., et al. 2021, AJ, 161, 107
McNeill, A., Gowanlock, M., Mommert, M., et al. 2023, AJ, 166, 152
Navarro-Meza, S., Aadland, E., & Trilling, D. 2021, RNAAS, 5, 111
Oelkers, R. J., Rodriguez, J. E., Stassun, K. G., et al. 2017, AJ, 155, 39
Pál, A., Szakáts, R., Kiss, C., et al. 2020, ApJS, 247, 26
pandas development team, T. 2020, pandas-dev/pandas: Pandas, latest, Zenodo, doi:10.5281/zenodo.3509134
Racca, G. D., Laureijs, R., Stagnaro, L., et al. 2016, Proc. SPIE, 9904, 99040O
Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, JATIS, 1, 014003
Robitaille, T. P., Tollerud, E. J., Greenfield, P., et al. 2013, A&A, 558, A33
Saha, A., Wang, Z., Matheson, T., et al. 2016, Proc. SPIE, 9910, 99100F
Scargle, J. D. 1982, ApJ, 263, 835
Sergeyev, A., Carry, B., Eggl, S., et al. 2023, LPICo, 2851, 2175
Süveges, M., Guy, L. P., Eyer, L., et al. 2015, MNRAS, 450, 2052
Tonry, J. L., Denneau, L., Heinze, A. N., et al. 2018, PASP, 130, 064505
Trilling, D. E., Gerdes, D. W., Juric, M., et al. 2023a, arXiv:2309.03417
Trilling, D. E., Gowanlock, M., Kramer, D., et al. 2023b, AJ, 165, 111
VanderPlas, J. T. 2018, ApJS, 236, 16
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, NatMe, 17, 261
Warner, B. D., Harris, A. W., & Pravec, P. 2009, Icar, 202, 134