

Automated eccentricity measurement from raw eclipsing binary light curves with intrinsic variability

L. W. IJspeert¹, A. Tkachenko¹, C. Johnston^{2,3,1}, A. Prša⁴, M. A. Wells⁴, and C. Aerts^{1,2,5}

¹ Institute of Astronomy, KU Leuven, Celestijnenlaan 200D, 3001 Leuven, Belgium
e-mail: luc.ijspeert@kuleuven.be

² Department of Astrophysics, IMAPP, Radboud University Nijmegen, P. O. Box 9010, 6500 GL Nijmegen, the Netherlands

³ Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Straße 1, 85741 Garching bei München, Germany

⁴ Villanova University, Department of Astrophysics and Planetary Sciences, 800 East Lancaster Avenue, Villanova, PA 19085, USA

⁵ Max Planck Institute for Astronomy, Königstuhl 17, 69117 Heidelberg, Germany

Received [month] dd, 2023; accepted [month] dd, 2023

ABSTRACT

Context. Eclipsing binary systems provide the opportunity to measure fundamental parameters of their component stars in a stellar-model independent way. This makes them ideal candidates for testing and calibrating theories of stellar structure and (tidal) evolution. Large photometric (space) surveys of today provide a wealth of data for both the discovery and the analysis of these systems. Even without spectroscopic follow up there is enough information in their photometric time series to warrant analysis, certainly if there is an added value present in the form of intrinsic variability like pulsations.

Aims. Our goal is to build a framework and its implementation for the homogeneous analysis of large numbers of eclipsing binary light curves aimed at the numerous high duty-cycle observations from space missions like TESS. The aim of this framework is to be quick and simple to run and to limit the user's time investment in obtaining, amongst other parameters, orbital eccentricities.

Methods. We develop a new and fully automated methodology for the analysis of eclipsing binary light curves with or without additional intrinsic variability. Our method includes a fast iterative prewhitening procedure resulting in a list of extracted sinusoids that is broadly applicable for purposes other than eclipses. After eclipses are identified and measured, orbital and stellar parameters are measured under the assumption of spherical stars of uniform brightness.

Results. We test our methodology in two settings: a set of synthetic light curves with known input and the catalogue of *Kepler* eclipsing binaries. The synthetic tests show that we can reliably recover the frequencies and amplitudes of the sinusoids included in the signal as well as the input binary parameters, albeit to varying degrees of accuracy. Recovery of the tangential component of eccentricity is most accurate and precise. *Kepler* results confirm a robust determination of orbital periods, with 81.8% of periods matching the catalogued ones. We present the eccentricities for this analysis and show that they broadly follow the theoretically expected pattern as function of the orbital period.

Conclusions. Our analysis methodology is shown to be capable of analysing large numbers of eclipsing binary light curves with no user intervention, and provide in that a basis for the further in-depth analysis of systems of particular interest as well as for statistical analysis at the sample level. Furthermore the computational performance of the frequency analysis, extracting hundreds of sinusoids from *Kepler* light curves in a few hours, demonstrates its value as a tool for a field like asteroseismology.

Key words. asteroseismology – binaries: eclipsing – ephemerides – methods: data analysis

1. Introduction

Eclipsing binary (EB) systems are an important source of physical properties of the stars (Torres et al. 2010). By being locked in a gravitational dance they provide the opportunity to directly measure their individual stellar masses, and by eclipsing each other they provide the opportunity to measure their radii (see Southworth 2012, for a brief but clear introduction to the topic). To achieve these measurements the right observational data is needed: a photometric time series of the brightness of the EB and enough data points of the radial velocities of both stars obtained through spectroscopy. Due to the nature of the observations, obtaining spectra is a lot more time-intensive than photometry, although both types of observations are seeing big boosts in the available amount of data: e.g. the CoRoT (Auvergne et al. 2009), *Kepler* (Koch et al. 2010) and TESS missions (Ricker

et al. 2015) for photometry and the Gaia-ESO (Gilmore et al. 2012), GALAH (De Silva et al. 2015), SDSS-IV APOGEE-2 (Blanton et al. 2017) and forthcoming SDSS-V MWM (Milky Way Mapper) (Kollmeier et al. 2017) surveys for spectroscopy.

Moreover, large photometric surveys are an ideal source for the discovery of variable stars like EBs (see e.g. Prša et al. 2011, 2022; Armstrong et al. 2016; IJspeert et al. 2021; Howard et al. 2022; Rowan et al. 2022). This means that there are large numbers of EBs with available photometric time series data that lack the multiple epochs of spectroscopic data needed to construct radial velocity curves. However, while less powerful, photometry alone can still provide valuable information about the physical properties of the system. To start, the orbital period determined from the eclipses can inform spectroscopic follow-up and it is an important ingredient for the further analysis of the system as well. Orbital eccentricities give the shape of the orbit and provide

a window into the stellar properties by being tied to dissipation within the stars. The argument of periastron and inclination define two of the three angles of the orientation of the orbit¹. The sum of radii scaled by the semi-major axis and ratio of radii provide sizes of the stars relative to the size of the orbit, which also ties back to dissipative processes. The ratio of surface brightness is related to the ratio of effective temperatures and when combined with the scaled radii tells us whether the two stars are more like twins or like far cousins.

There are several more properties that influence the shape of the light curve, one example being limb darkening, but these are usually more subtle effects. Furthermore, many stars, including those in binaries, are found to exhibit intrinsic variability like pulsations (Lampens 2021, and citations therein). Systems showing both phenomena, eclipses and pulsations, should not be cast aside, as they offer the biggest opportunity to learn by virtue of the information encoded in their light variations (Huber 2015). Indeed they are not always dismissed, as shown by the in-depth analyses by e.g. Johnston et al. (2019b, 2021); Maxted et al. (2020); Sekaran et al. (2020, 2021); Guo et al. (2022); Southworth & Bowman (2022); Southworth & Van Reeth (2022); Johnston et al. (2023). However, all of the named studies have in common that only one or a few, up to several tens, systems were treated. In their analysis, these systems often provide a complication: the amplitudes of for example pulsations can be comparable to or even exceeding the depth of eclipses in the light curve, rendering it impossible to only treat one of the sources of light variation separately. Any generalised approach should therefore account for both the eclipses present as well as any intrinsic variability that one or both stars exhibit.

To this day, many interesting eclipsing systems have been investigated manually in depth to great effect, providing physical insights on a case-by-case basis (e.g. Johnston et al. 2019a; Jennings et al. 2023; Pavlovski et al. 2023; Rosu et al. 2023, to name just a few). It is through these kinds of studies that we can calibrate in detail our models of stars, for example with regards to the discrepancy between the dynamical and evolutionary masses of intermediate- to high mass stars (e.g. Burkholder et al. 1997; Guinan et al. 2000; Weidner & Vink 2010; Massey et al. 2012; Tkachenko et al. 2020; Johnston 2021). This underlines the value of these systems for the calibration of our stellar models. With the number of eclipsing binaries in the thousands and increasing, investigating every one of them by hand in a finite amount of time is infeasible. If a portion of the analysis could be done in an automated way, this would give the researcher more time to take a more detailed look at a select number of promising systems. Additionally, there is a lot of power in numbers, as certain aspects of stellar physics can only be constrained by comparing with the population as a whole (see e.g. Van Eylen et al. 2016; Lurie et al. 2017; Wells & Prša 2021). Altogether, high levels of automation in the data analysis are needed to achieve many of the goals in stellar astrophysical research today.

Notable work towards automated pipelines for EB analysis includes DEBiL (Devor 2005), EBAS (Tamuz et al. 2006) and EBAl Prša et al. (2008). The DEBiL analysis code is part of a larger EB detection pipeline and uses simulated annealing and the downhill simplex method (Nelder & Mead 1965) to solve the inverse problem. EBAS has a strategy for making an initial guess of the parameters using grid searches, gradient descent and light curve geometry and then employs simulated annealing as well. DEBiL uses its own light curve model implementation of limb-

darkened spherical stars, whereas EBAS uses the more complex EBOP light curve generator. With EBAl the authors take a different approach and use machine learning with neural networks trained on the rigorous Wilson & Devinney (Wilson & Devinney 1971) model to arrive at a fast estimator for EB parameters. All three were applied to the then recently published database of OGLE (Udalski et al. 1992) EBs (Wyrzykowski et al. 2003), which at the time typically had on the order of a few hundred data points per target. However, while two of the mentioned pipelines include a measure of ellipsoidal variability, none of these automated analysis schemes account for intrinsic variability in the form of pulsations that we are specifically pursuing for their high astrophysical value.

We present a new and fully automated method and its implementation STAR SHADOW for the analysis of EB light curves with and without intrinsic variability. Starting from an input light curve the algorithm can provide orbital period, several orbit- and eclipse-specific parameters like eccentricity and a set of sinusoids which are added to the model of eclipses to fully describe the light variations in the time series. If further investigation into a systems eclipses or the remaining variability is desired, the removal of the other signal, be it the eclipses or the pulsations, is greatly simplified. We validate our method based on artificial data and apply it to a population of EBs from the *Kepler* database (Prša et al. 2011; Slawson et al. 2011; Kirk et al. 2016). to further validate the computed orbital periods. For the latter population, we deliver the orbital configurations, in particular eccentricities, and the properties of the variability intrinsic to the binary components. We proceed with a detailed description of the method and its application to a number of synthetic test cases.

2. Light curve analysis

We developed an automated recipe for the analysis of eclipsing binary light curves to handle large numbers of targets in a consistent, time efficient way and without manual involvement. The implementation of the recipe, called STAR SHADOW (Satellite Time series Analysis Routine using Sinusoids and Harmonics Automatedly for Double stars with Occultations and Waves), is publicly available on GitHub² and comes with a short guide to get started. Since one of the main aims is the application to data from the TESS mission, the algorithm works best for high duty-cycle satellite time series photometry and there are a few TESS-specific features like loading in TESS data products and recognition of TESS observing sectors. The algorithm is otherwise fully agnostic to the source of input light curve, so in the following text if we mention ‘sectors’, this can be taken to mean any subdivision of the data in the time domain. The two main goals are the determination of the orbital eccentricity and a measure for the significance of variability in the light curve after subtraction of the eclipses.

2.1. Top level overview

We provide a visual representation of the flow of the analysis steps in Figure 1.

The recipe starts with an iterative prewhitening procedure meant to capture all the statistically relevant (with regards to the Bayesian information criterion) variability in the light curve in a mathematical model of sine waves. The parameters of the sinusoids are optimised and, if not given beforehand, the orbital period is determined. Frequencies in the previously extracted list

¹ The third one being the longitude of the ascending node, which is generally not constrained.

² https://github.com/LucIJspeert/star_shadow

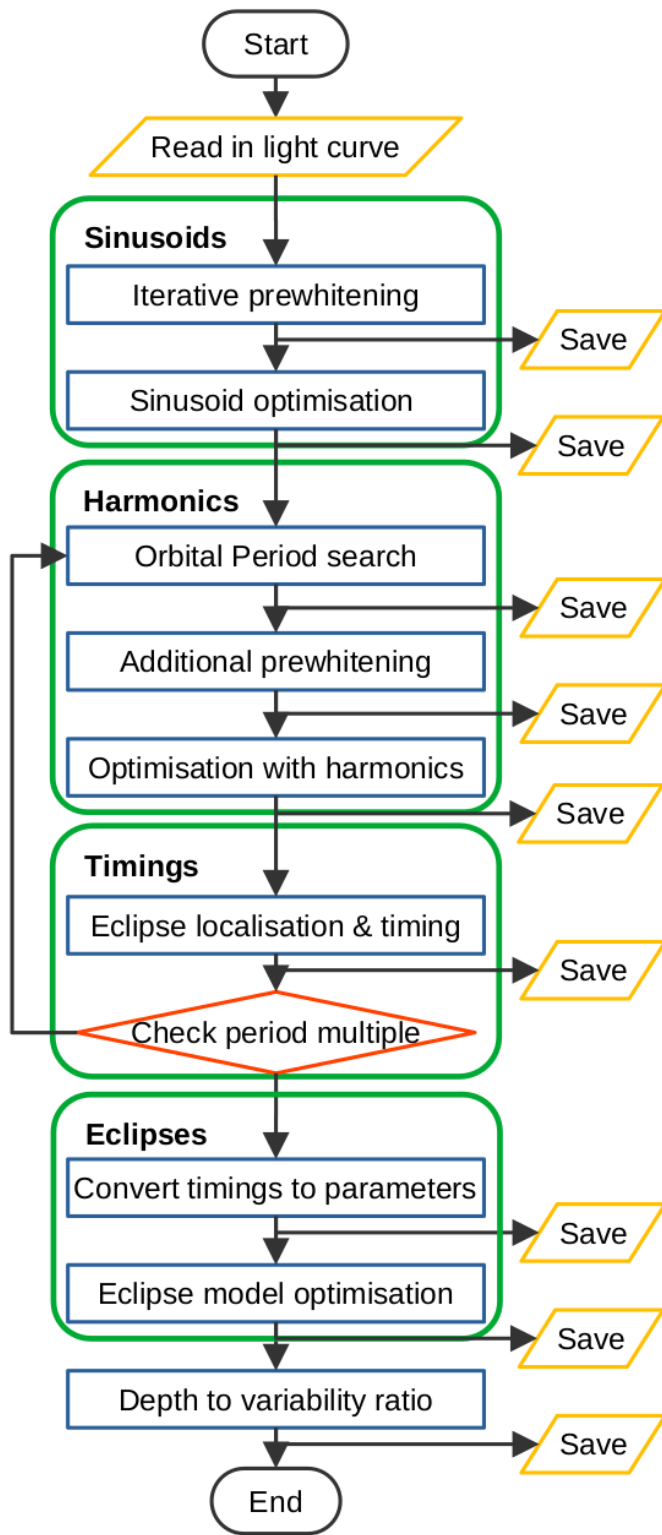


Fig. 1. Flowchart of the STAR SHADOW analysis steps. Blue rectangles contain analysis steps that are themselves grouped into four more general functionalities. Read and write operations are in yellow parallelograms, and the orange diamond denotes a decision point to redo a number of steps concerning the orbital period.

that match a theoretical orbital harmonic within a tolerance are forced to the exact multiples of the orbital frequency, thus replacing all frequencies within the tolerance by a single harmonic. If frequencies are closely spaced around the harmonics, this ap-

proach will remove too many frequencies from the model: an additional round of iterative prewhitening ensures that the model of the light curve is complete. The model is again optimised, now including the orbital period as free parameter, although its value tends to only change by a fraction of the period error estimate at this point, if at all.

The eclipses are located in the time series using their signatures in the first and second time derivatives, using only the model of harmonic sinusoids for the time derivatives to reduce the chance of mistaking variability for eclipses. Examples of what is meant by such eclipse signatures can be seen in Figure 2. The eclipse timings and depths are measured by finding specific extrema and zero points in the two time derivatives.

The eclipse timings and depths are translated into some of the orbital and physical parameters of the system: eccentricity, argument of periastron, inclination, sum of scaled radii, ratio of radii and ratio of surface brightness. This is done using Kepler’s second law and by iteratively solving a set of equations relating specific points in the orbit to the system’s properties. See further on in section 2.2, about ‘timings’, for more details.

The orbital and physical parameters are used as a starting point to fit a simple light curve model of eclipses to the data, to potentially increase the accuracy in some of the parameters. The model assumes spherical stars of uniform surface brightness and we refer to it later as just the ‘eclipse model’ for brevity. The fit is done in three steps: fit only the eclipse model to the original light curve, subtract the eclipse model and iteratively prewhiten the residuals, and then optimise the full model of eclipses and sinusoids together.

Finally, the ratios of eclipse depth to variability amplitude are computed to have a quick statistic of the balance between the two signals. At every step that adds, removes, or otherwise alters sinusoids, the extracted sinusoids are checked for significance against a signal-to-noise threshold, but not removed based on this check.

2.2. Detailed description

The model of the light curve constructed during the iterative prewhitening procedure also includes a linear trend for each sector. For TESS data, the linear trend is split up according to the first and second half of each sector, each piece being described by an independent slope and a y-intercept. This captures the average level of the light curve as well as remaining linear instrumental trends in the TESS data products, that may include sudden jumps in the middle of a sector due to momentum dumps of the gyroscopes. Where the the linear trend is split up can be defined by the user. The slope and y-intercept are calculated through their maximum likelihood estimators, and updated at each step of the subsequent iterations of the prewhitening as well as during optimisation procedures. See Appendix A for the formulae of these estimators.

The time axis of the data is mean-subtracted throughout the analysis, which ensures minimal correlation between parameters such as slope and y-intercept, and between frequency and phase. This means that sinusoid phases and eclipse times are measured with respect to this reference time. For the linear trends the mean is computed for each separate sector, so their y-intercepts are with respect to those time points.

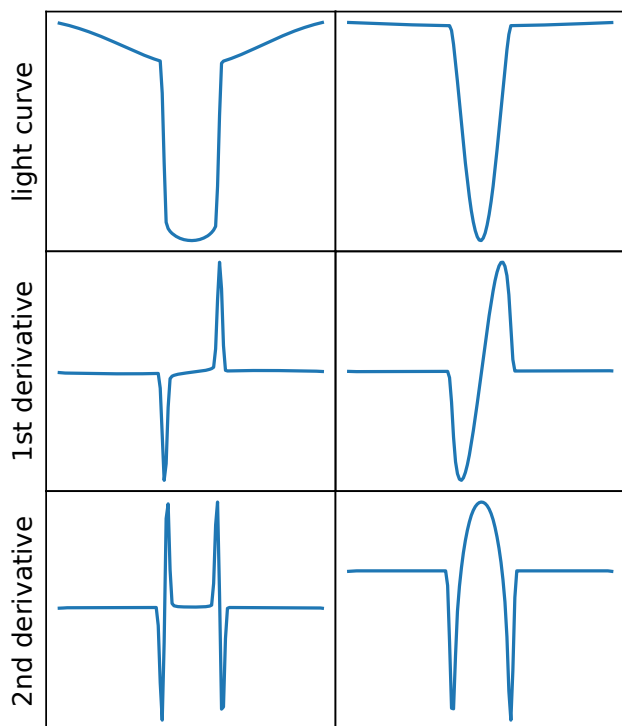


Fig. 2. Typical eclipse signatures in the derivatives of the light curve. The left column shows a flat-bottomed eclipse resulting in two separated features in the derivatives, while the V-shaped eclipse on the right leads to peaks merged into one feature. Adapted with permission from Ijspeert et al. (2021)

Frequency analysis

This Subsection encompasses the ‘Sinusoids’ and ‘Harmonics’ blocks in the flow diagram (Figure 1). In the iterative prewhitening procedure we build a model of the light curve in terms of sine waves. In this procedure the next sinusoid to be subtracted is picked by finding the highest amplitude peak in the Lomb-Scargle periodogram (Lomb 1976; Scargle 1982). The peak in the periodogram is oversampled³ by a factor 100 to obtain a precise position (i.e. frequency) and amplitude. The phase of the sinusoid is determined directly from the periodogram as well, using the formula described by Hocke (1998). A sinusoid is only accepted if the Bayesian information criterion (BIC) of the residuals is reduced by at least two, indicating at least positive evidence of its presence by rule of thumb (Raftery 1995; Kass & Raftery 1995). The BIC decreases with more white-noise like residuals (i.e. containing less information), but increases with additional free parameters of which each sinusoid initially has three (See Appendix A for a brief derivation). If the next sinusoid does not meet the requirement, the iteration is stopped. In the case of red noise, and depending on its power spectral density, high frequency peaks that are below the red noise amplitude but would be at high enough signal to noise are missed.

If at any point during the iterative prewhitening a frequency is found within the Rayleigh criterion of a previously extracted frequency, their peaks in the periodogram will have been blended together prior to subtraction. The Rayleigh criterion is calculated as $1.5/T$, where T stands for the total time base of observations. This blending means that the top of the true peak in the peri-

odogram has shifted and the extracted frequency is in a different location. A mechanism is introduced to counteract this shift, which checks each frequency that is being extracted for closely spaced frequencies in the previously extracted list. If a chain of closely spaced frequencies is found around the currently extracted frequency, all frequencies in the chain are iteratively updated by removing one of them at a time and re-extracting it from the updated periodogram around that frequency. Here ‘chain’ is used to mean any consecutive set of frequencies that are closer than the Rayleigh criterion to their direct neighbour. This extends to frequencies that are further than one Rayleigh criterion of the currently extracted frequency. The iteration of this sub-loop is terminated when the BIC of the residuals no longer decreases.

The sinusoid parameters can now have shifted with respect to their initial extraction and we may have to remove some sinusoids again based on the BIC. This is done in two parts, starting with removing individual sinusoids one by one and checking whether it results in a decrease in the BIC (of more than two). Secondly, it is checked whether chains of closely spaced frequencies (or any consecutive sub-chains) can be replaced by a single frequency and yield an improvement.

At the end of each step that involves extracting or updating sinusoid parameters, their significance is checked based on several criteria and stored alongside the parameters. This has no further effect on the analysis process. Sinusoids with an amplitude that is consistent with zero within three sigma are marked insignificant, and for frequencies that fall within three sigma of each other, the lowest amplitude sine wave is marked insignificant. The frequency dependent periodogram noise level is computed by taking a running average over the Lomb-Scargle periodogram of the residuals of the light curve using a frequency window of one per day. The signal-to-noise for each sinusoid is the amplitude divided by this periodogram noise level and those not meeting the threshold are marked insignificant. We use a signal-to-noise threshold as presented in Baran & Koen (2021, Equation 6), which is a function of the number of data points and computed for a false alarm probability of 0.1%. Candidate harmonics are selected as well, marking only those sinusoids that are within three sigma of a theoretical harmonic frequency (this selection is only relevant later in the analysis).

At the end of the frequency extraction it is necessary to do a multi-sinusoid non-linear optimisation to get closer still to the real frequencies present in the light curve (Bowman & Michielsen 2021, Sect. 3). We implement the `minimize` function from the Scipy (Virtanen et al. 2020) optimisation module and use the L-BFGS-B method with analytical gradients, see Appendix B for more insight into the choice of method and the mentioned gradient formulae. Doing such an optimisation whereby all sinusoids are simultaneously fitted is extremely time consuming, therefore to limit the time taken we fit in groups of 20 to 25 sinusoids. This reduces the fit performance in terms of the final BIC value reached by about one quarter of what it could have reached with a fully simultaneous fit in small scale tests, but the time saved is highly disproportional to this performance decrease (orders of magnitude). Groups are based on the amplitude of the sinusoids, bundling those with similar amplitudes together with the idea that those will have a similar level of influence on each other and on the likelihood function. The division is made by sorting the amplitudes in descending order and making the cut where the largest jump in amplitude occurs within the given constraints of minimum and maximum group size. Group size is not fixed to one number in order

³ Over the default $1/(10 \cdot \text{timespan})$ frequency sampling.

to increase amplitude disparity between groups. Frequencies, amplitudes and phases of the sinusoids in each group are fitted as free parameters in turn until all sinusoids have been optimised. Slopes and y-intercepts of all linear trends always remain free parameters throughout the fit.

Frequency analysis with harmonics

This Subsection encompasses the ‘Harmonics’ block in the flow diagram (Figure 1). We now introduce the orbital period as a free parameter.⁴ If known, an initial period can be provided, which is then further optimised. Alternatively, a period finding algorithm is used to obtain one.

The period finding algorithm makes use of the already obtained information of sinusoids by only searching at the extracted frequencies and fractions of those. A combined phase dispersion measure (Stellingwerf 1978) and Lomb-Scargle amplitude statistic is computed as described in Saha & Vivas (2017). In addition, this is multiplied with two more statistics that use the known list of frequencies and are geared towards improving the odds of finding the correct eclipsing binary orbital period. These two statistics are calculated as follows: 1) The number of harmonic sinusoids present in the extracted frequencies. At each tested orbital frequency, the theoretical orbital harmonic frequencies are computed and compared to the list of extracted frequencies. Those extracted frequencies that fall within the frequency resolution of a theoretical harmonic are selected and count towards the length of the harmonic series for that orbital frequency. 2) The filling factor of the harmonic series. The second statistic is also based on the harmonic series and captures the filling factor calculated as the number of harmonics present in the extracted frequency list divided by the total number of theoretically possible harmonics below the Nyquist frequency. The Nyquist frequency is defined as the sampling rate divided by two, where we use one over the minimum time step for the sampling rate.

After multiplying these four statistics the highest scoring frequency is taken as the initial orbital frequency globally speaking. The best orbital frequency is further refined locally by minimising the sum of squared distances between the series of extracted harmonic frequencies and the theoretical harmonic frequencies in a grid with sampling of 0.001% over a 1% interval around the orbital frequency.

As a final check, the two statistics on the series of harmonics as described above are calculated for several multiples of the orbital period: factors one-half, two, three, four and five, capturing all commonly mistaken multiples. There are two conditions for changing the period multiple, concerning both the filling factor itself and the multiplication of the two statistics of filling factor and harmonic series length. If the multiplication of the two statistics results in more than 1.1 times the value in this statistic than for the original period and the filling factor is at least 0.85 times that of the original period, the final orbital period is set to the multiple achieving the best score. Alternatively, if both the mentioned values exceed 0.95 specifically for doubling the period, the period is doubled. The threshold values for this decision are empirically determined based on the tests described in Sections 3 and 4.

The error on the period is estimated by taking the simple linear regression (SLR) uncertainty where each observed eclipse

is represented by a single timestamp. This error is only a function of the number of observed eclipses and the uncertainty on the individual timestamps, which is taken to be half the integration time for the observations at this stage. The number of observed eclipses is approximated by rounding down the time base divided by the orbital period, thus assuming no gaps.

Once we have a binary period we select the orbital harmonic frequencies to couple them. All frequencies falling within the Rayleigh criterion of the theoretical harmonics are removed and replaced by exact multiples of the orbital frequency. The new amplitudes and phases are obtained from the periodogram. From this point the orbital harmonics will have a reduced number of free parameters, from three for unconstrained sinusoids to two, while adding one free parameter for the orbital period. Despite the reduction in parameters and updated harmonic sinusoids, it is possible for this step to increase the overall BIC of the residuals, as too many sinusoids may have been removed.

Since harmonics have one free parameter fewer, it is possible that additional harmonics can be extracted from the periodogram while decreasing the BIC of the residuals more than two. This is attempted by calculating the amplitude and phase at harmonic frequencies not yet occupied in the frequency list, up to the Nyquist frequency, and checking the significance criterion before they are added to the sinusoidal model of the light curve. Additionally, due to the removal of candidate harmonics it is likely that more unconstrained sinusoids can be extracted analogously to the initial prewhitening. After adding sinusoids to the list we check again whether individual frequencies may have to be removed, or that closely spaced chains of frequencies can be replaced by a single frequency based on the BIC of the residuals. If a chain of close frequencies contains a coupled harmonic sinusoid, it is attempted to replace the chain by the harmonic and the harmonic is not removed. The harmonic does not change in frequency but its amplitude and phase can change.

Finally, a second multi-sinusoid non-linear optimisation is performed analogously to the first. The difference is that there is now a set of harmonic sinusoids that is coupled to the orbital period. The period is also optimised along with the amplitudes and phases of the harmonics, which are not sub-divided into groups regardless of their number.

Eclipse timing measurements

This Subsection encompasses the ‘Timings’ block in the flow diagram (Figure 1). For the purpose of locating the eclipses, the eclipses are separated out from other variability in the light curve as much as possible by only using the model of orbital harmonic sinusoids (henceforth: harmonic model). Variability at the harmonic frequencies other than the eclipses cannot be separated yet, which can cause problems when that variability has high frequencies relative to the orbital frequency. For this reason, we initially determine the approximate eclipse positions using orbital harmonics up to and including the twentieth harmonic. If this fails, it is attempted again with the first forty harmonics and finally with all the orbital harmonics. We provide a visualisation in Figure 3 to aid in the explanation of analysis steps below, matching the markers in the Figure with numbers and letters between brackets in the text.

The first and second analytically computed time derivatives of the harmonic model inform the automated recognition of the eclipse signatures. The highest peaks in the absolute first derivative belong to the steepest slopes, which in most eclipsing binary

⁴ Up to this point, the prewhitening steps are applicable to any time series and could be used as a convenient and quick method on its own.




Fig. 3. Typical eclipse signatures in the derivatives of the light curve with markers pointing to important features for the automated eclipse detection and timing measurement. Note that points 7 in the left panel of first derivatives are not truly local minima, but marked as such for illustrative purpose. Adapted with permission from IJspeert et al. (2021)

light curves will mean an eclipse ingress or egress. The sign of the first derivative at these locations gives the direction of the slope, so that we can distinguish the start of the eclipse from the end of the eclipse. We proceed by determining the first and last contact points⁵ of the eclipses by following several steps, as described below. Starting from the extrema in the first derivative (1) and moving away from the centre of the eclipse, we find its first zero point (2). Between the extrema and these zero points, there is a minimum in the second derivative (3) where the curvature of the eclipse is greatest (most negative). The minima in the second derivative are taken to be the inner limits for the position of the first and last contact points. We check for the presence of local minima in the absolute first derivative (4) between the zero points in the first derivative and the minima in the second derivative. If present, these local minima replace the zero points in the first derivative as the outer limit for the position of the first and last contact points. The times of first and last contact (a) are taken to be midway between the described inner and outer limits, and the interval (3-4) is taken to be the three sigma error region.

Towards the centre of the eclipse we determine the first and last internal tangency points⁶, analogously to the contact points albeit in opposite direction. Starting again from the extrema in the first derivative (1) and moving towards the centre of the

⁵ Contact refers to the point in time when the observer sees the projected discs of the two stars touch while they are not (yet) covering each other. In other words the start and end of the eclipse. This always happens per definition of an eclipse.

⁶ Internal tangency refers to the point in time when the observer sees the projected discs of the two stars touch tangentially while one is in front of the other. This does not necessarily happen in every eclipse.

eclipse, we find its first inner zero point (5). Between the extrema and these zero points, there is a maximum in the second derivative (6) where the curvature of the eclipse is greatest (most positive). The maxima in the second derivative are taken to be the outer limits for the position of the first and last tangency points. We check for the presence of local minima in the absolute first derivative (7) between the zero points in the first derivative and the maxima in the second derivative. If present, these local minima replace the inner zero points in the first derivative as the inner limit for the position of the first and last tangency points. If we find two distinct maxima in the second derivative and the second derivative reaches zero in between them, the first and last internal tangency points (b) are defined as the midway between their respective limits and the interval (6-7) is taken to be the three sigma error region. The internal tangency points now delimit the flattened off bottom of the eclipse. If instead we find that these two maxima coincide, or the curvature does not reach zero in between them, no such flat bottom is defined.

The eclipse minima (c) are defined to be in the middle between the internal tangency points, and not as the actual minimum in the harmonic model. This is more robust in any cases with variability at harmonic frequencies that does not belong to the eclipses. Possible asymmetry can still be accounted for because the internal tangency points may be anywhere between the contact points.

The best consecutive primary and secondary eclipse combination is chosen by taking those with the largest combined depth. If no secondary is found, or no eclipses are found at all, the analysis is stopped. The method for determination of the orbital eccentricity does not hold for systems with no detected secondary eclipse. This method of locating the eclipses was first described in IJspeert et al. (2021, Section 3), but was adapted and simplified for the use with analytic sums of sine waves instead of raw data. Note that in the left panel of first derivatives in Figure 3, points 7 are not truly local minima in the portrayed derivative of the theoretical eclipse model. Using finite sums of sine waves generally results in many more local minima and maxima than is the case for a theoretical model of eclipses, because of the inflection points around (sharp) bends, thus allowing for the simple identification of points 7 as local minima in the Figure in case a flat eclipse bottom is present.

Once two eclipses have been detected their timing measurements are refined using a larger number of harmonics if possible. Limiting to lower harmonics works to make detection more robust, but the downside is that eclipses look wider and shallower. An altered version of the eclipse detection algorithm is used that has the benefit of starting off from the previously detected positions and is more robust against high frequency interference. The characteristic points measured are the same, but the method of arriving at them is different.

Starting off from previously determined inner zero points in the first derivative (5), we adjust them to the different number of sinusoids included in the first derivative by following the slope to zero. Subsequently, updated extrema in the first derivative (1) are determined by searching between the inner zero points and the previous extrema in the first derivative. Moving away from the eclipse centre from the newly found extrema in the first derivative, we find the outer zero points (2). If secondary (local) extrema occur within the outer zero points, and these are sufficiently high (> 80% of the height), the extrema in the first derivative (1) are shifted to those locations. The minima in the second derivative (3) are now found between the extrema and outer zero points in the first derivative. Local minima in the first derivative (4) are found by moving outward from the minima in

the second derivative. The inner zero points (5) are now adjusted again, by moving inward to zero from the extrema in the first derivative. The maxima in the second derivative (6) are found between the extrema and inner zero points in the first derivative. Inner local minima in the first derivative (7) are obtained by moving inward along the first derivative curve from the maxima in the second derivative. If the ingresses or egresses are found to be less than a quarter the width of initial localisations suggest (obtained with fewer harmonics), it may indicate that we do not capture the full extent of the eclipse due to variability not related to the eclipses. We adjust the outer eclipse points (1,2,3,4) further outward by moving to the next extrema in the first derivative (1) and repeating the steps to find the other points. This adjustment is only accepted if it does not cross the original outer edges (4) as determined in the localisation step, and if it increases the depth by more than twenty percent.

Improved error estimates for the individual timings and error estimates for the individual depths are determined using the noise level and the slope of the eclipse in- and egress. The noise level is calculated as the standard deviation of the residuals of the light curve. For errors in the individual eclipse times of ingress and egress we determine an average slope of each ingress and egress by fitting a straight line. Dividing the noise level by this slope gives the time it takes the flux level to cross the noise level, thus providing an estimate of the timing precision for contact points. Errors in the individual eclipse times are derived from this by taking half and adding in square to the error estimates obtained during eclipse detection:

$$\sigma_{t,in/egress} = \sqrt{\frac{1}{4}\sigma_{t,noise/slope}^2 + \sigma_{t,detection}^2} \quad (1)$$

Analogous to the orbital period error the error on the final eclipse timings, over all observed eclipses, is determined from the SLR uncertainty, which is only a function of the number of observed eclipses and the uncertainty on the individual timings given by Equation 1. The error for the orbital period is also recalculated here, now using the averaged error on the primary and secondary eclipse times as input instead of the time stamp error.

The error in the flux level at first and last contact is estimated as being the noise level, and the error at the bottom of the eclipse is either the noise level (if there is no flat bottom), or the standard deviation of the flux measurements from first to last internal tangency (in case there is a flat bottom). These flux level errors then translate into depth errors by adding them in square:

$$\sigma_{depth} = \sqrt{\frac{1}{4}\sigma_{first\ contact}^2 + \frac{1}{4}\sigma_{last\ contact}^2 + \sigma_{bottom}^2} \quad (2)$$

the factor one-fourth comes from averaging the flux levels for the first and last contact points in the calculation for the eclipse depth. To obtain the error on the final eclipse depths, over all observed eclipses, we divide Equation 2 by the square root of the number of eclipses.

Eclipses are rejected under certain conditions, at the end of the detection phase but also at points during the determination of the best candidates for primary and secondary eclipse. Most notably we apply a significance criterion to the depth and duration of the eclipses based on the estimated errors and a signal-to-noise threshold. The depth must exceed one sigma of the estimated error and one half the standard deviation of the residuals whereas the duration must exceed three sigma of the estimated error. Examples of other criteria, applied during selection of candidate eclipses, look at differences between

overlapping candidates and whether a candidate greatly changes in depth between harmonic models with all or only low frequency harmonics (indicating that it is more likely unrelated signal). These selection criteria are meant to filter out as many false positive candidates before choosing the final set of eclipses to use in the rest of the analysis. However, these criteria or the conditions applied at the end may result in only a primary eclipse (or none), in which case this is logged and the analysis ends.

Eclipse analysis

This subsection encompasses the ‘Eclipses’ block in the flow diagram (Figure 1). Determination of orbital and physical properties including eccentricity, argument of periastron, orbital inclination, sum of scaled radii, ratio of radii and ratio of surface brightness is initially done through the use of formulae coupling these properties to the timings and depths of the primary and secondary eclipse. These formulae can be found in e.g. Kopal (1959), but some will be repeated here in altered form for completeness. As the main link between time and orbital elements, we need Kepler’s second law in integral form:

$$\frac{2\pi(t_b - t_a)}{P} = (1 - e^2)^{\frac{3}{2}} \int_{\nu_a}^{\nu_b} \frac{d\nu}{(1 + e \cos(\nu))^2}, \quad (3)$$

where t_a and t_b represent time points in the orbit, P is the orbital period, e the eccentricity and ν_a and ν_b the true anomaly angles corresponding to the time points. This equation has an analytic solution which we put in the Appendix for brevity (A.32). Since we do not know the true anomaly (ν) angles beforehand we need to break it down into argument of periastron (ω) and phase angle (θ):

$$\nu = \frac{\pi}{2} - \omega + \theta. \quad (4)$$

The argument of periastron is a free parameter defining part of the orbital orientation in the sky, the phase angle gives the position along the orbit and is zero near primary minimum and π near secondary minimum. We now only need to determine the phase angles at several strategic points in the orbit to calculate useful values with our integral (3). These points constitute the times of minima, times of contact and for certain cases times of internal tangency as well. Times of minima are reached when the projected orbital separation reaches a minimum, or when its derivative becomes zero:

$$0 = e \cos(\theta - \omega) + \sin^2(i) \cos(\theta) (\sin(\theta) - e \cos(\omega)), \quad (5)$$

with i the orbital inclination. From this formula we can see that in circular orbits the minima are at phase angles of exactly zero and π . Times of contact are reached when the projected distance between the two stars is equal to the sum of both radii (scaled by the semi-major axis), leading to the equation:

$$\sqrt{1 - \sin^2(i) \cos^2(\phi)} = \frac{r_1 + r_2}{a(1 - e^2)} (1 \pm e \sin(\omega \pm \phi)), \quad (6)$$

where we introduce the angles ϕ which are the phase angles of the contact points as measured from zero and from π for the primary and secondary eclipse, respectively. In their exact form, Equations 5 and 6 do not have analytical solutions. Approximate analytical solutions do exist, but these are only valid for low eccentricities.

To be able to obtain accurate results for all eccentricities, we use an iterative form of Kepler's second law and solve numerically for ψ :

$$\frac{2\pi(t_2 - t_1)}{P} = \psi - \sin(\psi), \quad (7)$$

where ψ relates to the component $e\cos(\omega)$ of the eccentricity. We can write:

$$e\cos(\omega) = \frac{\tan\left(\frac{\psi-\pi}{2}\right) \sqrt{1 - (e\sin(\omega))^2}}{\sqrt{1 + \tan^2\left(\frac{\psi-\pi}{2}\right)}}, \quad (8)$$

which depends on an estimate of $e\sin(\omega)$. For $e\sin(\omega)$ we can write a formula with explicit dependence on inclination, as well as auxiliary angle ϕ_0 , which further simplifies with the assumption of small ϕ_0 and $i = 90^\circ$:

$$e\sin(\omega) = \frac{\pi(\tau_{1,1} + \tau_{1,2} - \tau_{2,1} - \tau_{2,2})}{2P\sin(\phi_0)} \frac{\sin^2(i)\sin^2(\phi_0)}{1 - \sin^2(i)(1 + \sin^2(\phi_0))}, \quad (9)$$

$$\approx \frac{d_2 - d_1}{d_2 + d_1}, \quad (10)$$

introducing the notation τ as the time duration from conjunction to either contact point and the subscript $\tau_{n,m}$ is used to indicate primary ($n = 1$) or secondary ($n = 2$) eclipse and first ($m = 1$) or last ($m = 2$) contact. Eclipse durations are denoted with d_n here. We use formula 9 with $i = 90^\circ$ and the approximate formula for the auxiliary angle ϕ_0 as in Kopal (1959):

$$\phi_0 \approx \frac{\pi}{2P} (\tau_{1,1} + \tau_{1,2} + \tau_{2,1} + \tau_{2,2}). \quad (11)$$

Angle ϕ_0 relates to the sum of scaled radii in the following manner:

$$\sqrt{1 - \sin^2(i) \cos^2(\phi_0)} = \frac{r_1 + r_2}{a(1 - e^2)}, \quad (12)$$

with a the semi-major axis. After obtaining the components of the eccentricity, we use an iterative procedure to improve the value of ϕ_0 by computing the sum of eclipse durations using Kepler's second law (Equation A.32) and Equation 6. A fast global minimisation method is used to find inclination, ratio of radii and ratio of surface brightness by comparing measured eclipse times and depths against the values computed using Equations A.32, 5, 6 and 13.

All of the above assumes spherical stars of uniform brightness. Error estimates are made using a combination of analytical error formulae and the Monte Carlo approach of importance sampling. We go into more detail for this step and the error formulae in Appendix B.

Fitting eclipse models to a light curve with no prior information usually does not work as the parameter space is degenerate and the optimisation is prone to get stuck in local minima. Having obtained the parameters from the eclipse timings and depths in the previous step, we avoid this problem by being able to start close to the optimal solution.

We construct a simple physical eclipse model by assuming spherical stars of uniform surface brightness and calculating the fraction of light lost due to one star covering the other, which we will refer to as the eclipse model. We do not include light from

a third body in this model. Subtracting the fractional light loss from one gives us the leftover normalised flux l :

$$l = 1 - \frac{A_{covered}}{\pi r_1^2 + \pi r_2^2 \frac{sb_2}{sb_1}}, \quad (13)$$

where sb_1 and sb_2 stand for the surface brightness of star 1 and 2, respectively. This formula gives the normalised flux for the primary eclipse: we need to multiply the right hand fraction by $\frac{sb_2}{sb_1}$ to obtain the flux for the secondary eclipse. The covered area $A_{covered}$ is the result of two overlapping circles as a function of the distance between their centres s and is computed as follows:

$$A = r_1^2 \arccos\left(\frac{s^2 + r_1^2 - r_2^2}{2sr_1}\right) + r_2^2 \arccos\left(\frac{s^2 + r_2^2 - r_1^2}{2sr_2}\right) - r_1 r_2 \sqrt{1 - \left(\frac{r_1^2 + r_2^2 - s^2}{2r_1 r_2}\right)^2}. \quad (14)$$

The separation between the two centres can be expressed in terms of the orbital parameters and as a function of the phase angle:

$$s = \sqrt{\frac{(1 - e^2)^2 (1 - \sin^2(i) \cos^2(\theta))}{(1 - e \sin(\theta - \omega))^2}}. \quad (15)$$

In the three equations above the stellar radii r_i and the separation s are in units of the semi-major axis a , but this is left implicit for the simplicity of the formulae. The model is calculated at a regular interval of 0.001 rad in θ , which is converted into times using Kepler's law (Equation A.32). Linear interpolation is used to get the values at the times of the observations.

This eclipse model is first fit to the data separately, starting from the previously obtained parameters, after which it is subtracted from the light curve and the residuals are prewhitened analogously to the earlier steps. The obtained model of sinusoids, linear model and the eclipse model are then optimised simultaneously. We implement two approaches here: Markov Chain Monte Carlo (MCMC) sampling with PyMC3 (Salvatier et al. 2016) and the Scipy minimize with method L-BFGS-B as before (default option)⁷. Since MCMC can deal with high dimensionality, it gives the full correlation structure as an output and thus additionally provide statistical errors estimates. However, these estimates are found to be small compared to our other error estimates (and thus underestimated, see Section 3), and the procedure is slower than the fitting method. Note that the orbital period and time of primary minimum are fixed and no longer free parameters at this point.

2.3. Remaining variability

The level of variability compared to the eclipse depths is computed to get an insight in the strength of the presence of potential pulsations or the harmonic sinusoids that are not captured by the eclipse model. Standard deviations are computed of the residuals of four different models: a model with no sinusoids, a model with only the first and second harmonic sinusoids, a model with only non-harmonic sinusoids and the full model including all sinusoids. All four models include the linear trend

⁷ For the fitting method the sinusoids are again subdivided into groups.

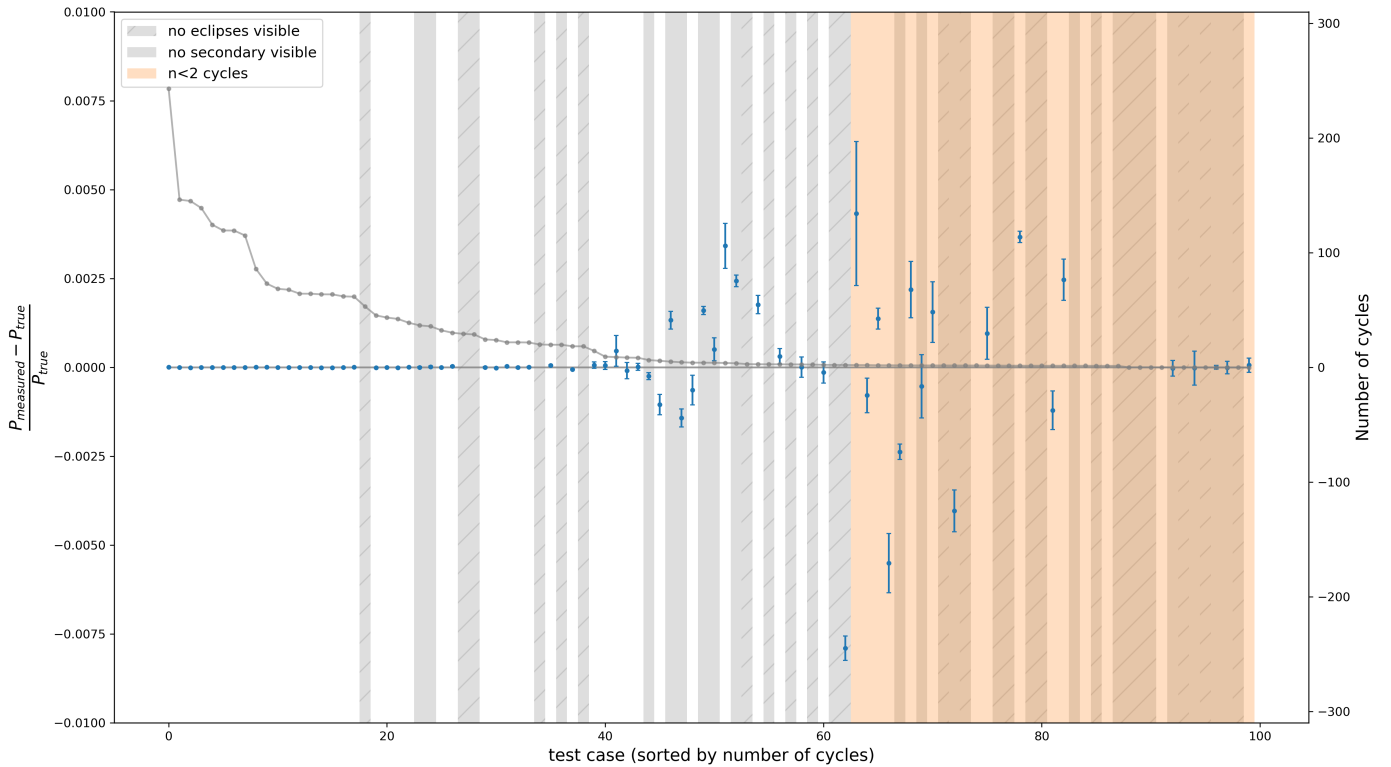


Fig. 4. Blue points with error bars: fractional difference between the measured and the true orbital period with simple linear regression error estimates, sorted by descending number of cycles. Orange points with error bars: same as blue but with MCMC errors. Grey line with points: number of cycles calculated as period divided by time base, plotted on the right axis. The period error scales strongly with the number of cycles. Note that the number of cycles for cases ~ 60 to 100 is close to one, and cases that are shaded grey and hatched do not show eclipses.

and the eclipse model. Then we take the ratio of primary and of secondary eclipse depth to each of these four standard deviations.

While we do not explore them in this work, additional steps of analysis on the remaining variability could be easily integrated into the framework.

3. Application to synthetic data

We produced a set of synthetic test light curves with randomly sampled parameters, including orbital period, eccentricity, argument of periastron, inclination, scaled radii, surface brightness ratio, third light and a number of pulsations represented by sinusoids. Light curves of a single TESS sector (27.4 days) and a full year (13 sectors) at 30-minute cadence were made to mimic some of the possible lengths of TESS time series. The eclipsing binary light curve models were produced with `ellc` (Maxted 2016) and sine waves were added to act as intrinsic variability of various frequencies and amplitudes, finally adding Gaussian noise to imitate real data. Due to the random nature of the generated parameters, some of the synthetic light curves contain no eclipses or have such shallow eclipses that they are impossible to spot by eye. These systems are marked in Figure 4 and later removed for a cleaner overview of the results. From the viewpoint of the analysis, there are several indicators for such difficult cases, most importantly the completion of all stages of the analysis. If the analysis was completed, the following parameters provide a handle on possible erroneous results: the number of eclipse cycles (period over time base) and depth of the secondary eclipse. We will go into some more detail on both of these below.




Fig. 5. Histogram and KDE of the orbital period χ -values, using the SLR uncertainties. The KDE is scaled to the height of the histogram.

The analysis algorithm was run with no prior knowledge of the orbital period, which is the most important parameter for the rest of the analysis to produce accurate results. We show the performance of the global period search in Figure 4. In all relevant

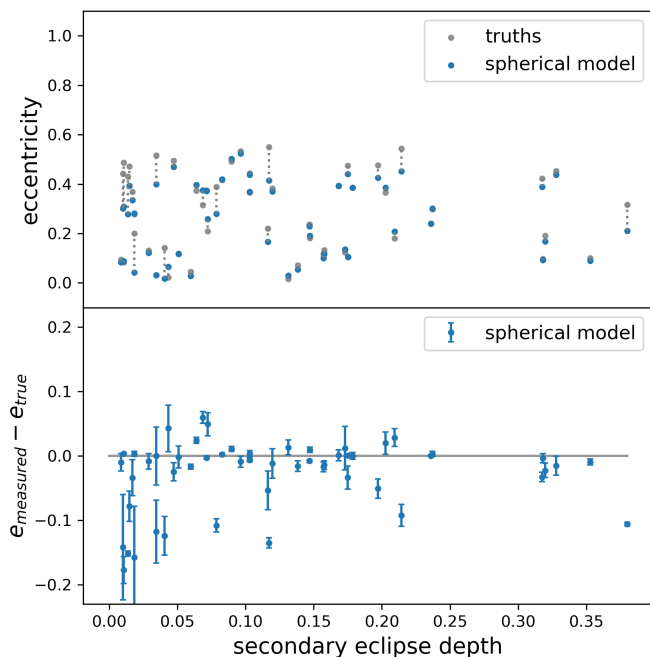


Fig. 6. Top panel: true (grey) and measured (coloured) eccentricities, bottom panel: deviation from the true eccentricity against the secondary eclipse depth. Scatter in the bottom panel is seen to increase to the left, with decreasing eclipse depth, as the eccentricity depends critically on the identification and measurement of the secondary eclipse.

cases ($n_{cycles} \geq 2$, visible eclipses) the algorithm was able to pick the correct multiple of the true period to better than 0.34% precision. More specifically, using the same selection criteria, the mean absolute deviation from the true period is 0.031%, with a median of 0.0009%. Further restricting to cases that have more than ten eclipse cycles reduces those numbers to a mean of 0.0011% and median of 0.00048%.

As Figure 4 is sorted by the number of cycles, we can see the importance of this parameter as an indicator of the accuracy in orbital period. This dependence on the number of cycles is taken into account in determining error values, and in the majority of cases the estimated SLR errors include the true period value within one or two sigma. Figure 5 shows the difference between the measured and true period divided by the SLR errors (χ_p) in a histogram and kernel density estimation (KDE), where the KDE is scaled such that its height is comparable to the histogram. The distribution has a broad peak at zero (width > 1), indicating accurate period measurements with estimates of their precision that are over all underestimating the true errors made.

Eccentricities

The further analysis depends on the secondary eclipse and thus its correct identification and the measurements of time of minimum, points of in- and egress and depth. This becomes more difficult with decreasing secondary eclipse depth, as evidenced by the increasing scatter to the left in Figure 6. True and measured eccentricity values are shown in the top panel for reference while the bottom panel shows the deviation from the truth. What actually causes the measurements to be more difficult for lower secondary depths is both the decrease in signal-to-noise as well as the ratio between the secondary eclipse depth and the ampli-

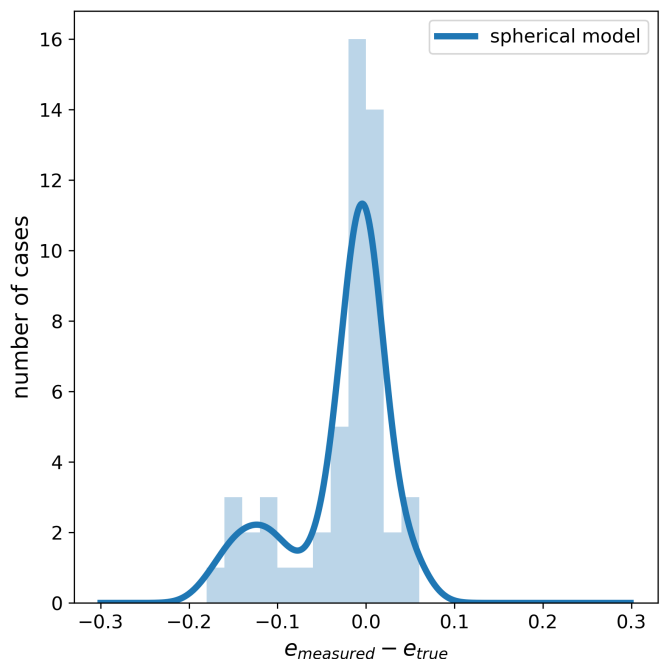


Fig. 7. Histogram and KDE of the differences between measured and true eccentricity. Next to the sharp peak at zero and within 0.1 deviation, there is a relatively small second maximum where eccentricities are underestimated between -0.2 and -0.1. The KDE is scaled to the height of the histogram.




Fig. 8. Histogram and KDE of the eccentricity χ -values, using error formula uncertainties. The bi-modality seen in the absolute differences has largely disappeared. The KDE is scaled to the height of the histogram.

tude of any leftover variability with high frequency compared to the orbital frequency. Both of these affect the precision with which we can localise the eclipse, and intrinsic variability can additionally impact the accuracy by either shifting the eclipse position or outright cause the method to miss the second eclipse entirely in the most extreme of cases. We refer to Figure C.26 in

the Appendix for the tangential components of the eccentricity. Specifically looking into the few cases with large (negative) eccentricity offsets at eclipse depths above 0.05 we see that the tangential components only constitute a fraction of the difference, up to about 0.03 in $\text{ecos}(\omega)$.

Deviations in eccentricity generally stay within 0.1 and although error estimates are representative in many cases, we also see cases with underestimated errors that are not explained by their position along the axis of secondary eclipse depth. We see a second maximum in the distribution of absolute differences between measured and true eccentricities (see Figure C.2), between values of -0.2 and -0.1, underestimating the true eccentricities. Inspecting these cases, we find no overarching indicator for their deviation, as their eclipses are correctly identified. We do find their deviations to stem nearly completely from the radial component, $e \sin(\omega)$. This is caused by inaccurate determination of eclipse times, that are in some cases easily seen in their analysis output but not in others (See Figures C.28 and C.30 for examples of this). Figure C.3 shows the distribution of the true deviations relative to the estimated errors, χ_e , which has a large width (> 1), meaning errors are underestimated over the whole. The bi-modality of the former distribution has largely disappeared, which indicates that the errors are sufficiently large for the second maximum to be incorporated into the main distribution of points. Both the eccentricities obtained from translating the timings (not shown here, see Appendix C) and those obtained from the eclipse model (shown here) have a similar spread and the mean for both distributions lies at zero, suggesting no significant bias towards lower or higher values.

The MCMC method also provides error values (not shown here), these are generally smaller than the previous estimates and more often than not missing the true values. This shows that while the MCMC sampling provides an overview of the full correlation structure it does not capture the full extent of the error; it only captures the statistical error made when fitting a model with no uncertainties to noisy data. Results obtained using the L-BFGS-B fitting method are qualitatively the same as those obtained through the MCMC sampling method.

Distributions for the deviations from truth of the parameters argument of periastron, inclination, sum of scaled radii, ratio of radii and ratio of surface brightness are given in Appendix C. We note a few key takeaway points from those Figures. There is a good agreement between the overall distributions of the parameters values from translating the timings and those obtained after the eclipse model optimisation step. Two exceptions seem to be the radial component of eccentricity, where the distribution for the eclipse model is shifted slightly towards overestimating this parameter by about two percent, and the auxiliary angle ϕ_0 , which is slightly underestimated by the eclipse model by less than a percent. The tangential component of the eccentricity is accurate and precise with a distribution width of about 0.01, but errors in a minority of cases are underestimated, resulting in long tails for its distribution of true deviations relative to the estimated errors. The sum of scaled radii shows good agreement with ground truth: its distribution is less than 0.05 in width, be it with a heavy tail on the high end. Orbital inclination shows a slightly wider spread, and no correlation is seen between inclination deviations and third light. The ratio of radii measurement performs worst over all: we attribute this to the fact that for grazing eclipses, and ignoring limb darkening, there is not enough information in the light curve to strongly constrain the individual stellar radii. We do note that the errors for ratio of radii appear overestimated for the majority of cases, narrowing down the central peak in its

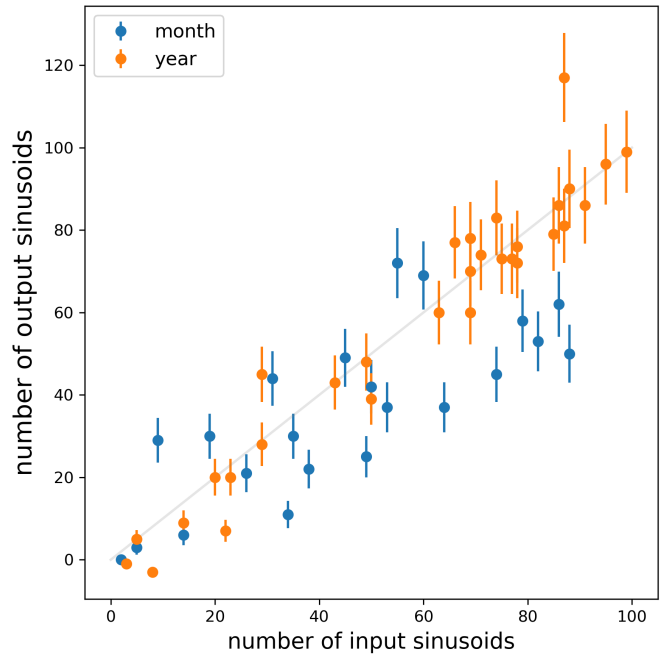


Fig. 9. The number of extracted independent sinusoids ($n - n_h$) versus the true number of sinusoids put into the light curve (n_{true}). Harmonic sinusoids are filtered out as they capture residuals related to the eclipse model that is unrelated to the number of input sinusoids. Points are divided up into month-long time series (blue) and year-long time series (orange) and the error bars are calculated as the square root of the number $n - n_h$.

distribution of deviations relative to the error estimates. The addition of limb darkening to our eclipse model may improve the predictive power in this area, but for this specific set of synthetic light curves, other variability would dominate over this effect in most cases. The surface brightness ratio performs better than the radii ratio on an absolute scale, but its error estimates are more underestimated and this gives it a poor result in the relative scale.

The number of sinusoids used in generating the synthetic light curves is randomly generated as well as their parameters. It is not evident a priori that we would recover this number: iterative prewhitening procedures tend to vary widely in the amount of sinusoids found based on stopping criteria and other aspects. Our approach is to rely on a statistical measure for the information held in the data (the BIC) during prewhitening, and only select based on significance criteria at the very end. Our final light curve model constitutes an eclipse model, linear trend and a sum of sinusoids. We now look only at those sinusoids in the final model that pass the significance criteria, which includes the signal-to-noise threshold. Most of the harmonics found in this filtered list contain the leftovers of the eclipse signal not captured by the eclipse model, like ellipsoidal variability and effects of limb darkening. Therefore it makes sense that we do not see a correlation with the number of harmonics in this list and the number of input sinusoids. We subtract the number of harmonics (n_h) from the total number of filtered sinusoids (n) to arrive at the number of independent frequencies ($n - n_h$), which is shown in Figure 9. We find that this number strongly correlates with the true number of frequencies, and that the results for longer time series tend to lie closer to the diagonal over all.

We take the square root of $n - n_h$ as error estimate, and compute the histogram and KDE for the deviations from the true




Fig. 10. Histogram and KDE of the number of sinusoids χ -values, where errors are taken to be the square root of the amount $(n - n_h)$. The histogram is split up into light curves of a month in length and those of a year in length: the shorter time base tends to give underestimates for the number of sinusoids present. The KDE is scaled to the height of the histogram.




Fig. 11. Histogram and KDE of the frequency χ -values for one case of the synthetic light curves, where errors are calculated with the standard formulae. The KDE is scaled to the height of the histogram.

value divided by the error to arrive at the distribution shown in Figure 10. When separated into cases with month-long and year-long light curves, the former have consistently slightly underestimated numbers of sinusoids while the latter have unbiased estimates centred at zero.

One case with a large number of added sinusoids is picked out to look further into the accuracy of the parameters of the extracted sinusoids. We compare the true frequency values with the measured ones by pairing the closest matches together. Figure 11 shows the accuracy and precision of measured frequencies, showing that the error formulae agree with the deviations from ground truth. The same is true for the amplitudes and to a lesser degree for the phase measurements, where the latter has a wider distribution by about 50%.

4. Kepler EB catalogue

For the purpose of further testing the period finding capability of STAR SHADOW against a vetted sample of real stars, we apply our methods to the 30-minute cadence light curves for the *Kepler* EB Catalog targets collected by Prša et al. (2011); Slawson et al. (2011) and Kirk et al. (2016).⁸ In addition we present the results for the eccentricities found for this data set. Of the 2920 detrended light curves in the catalogue, 2834 completed analysis at least up to and including the period measurement and of those 2001 got an eccentricity measurement (over all morphologies).

Generally we see a high level of agreement between the periods found by our method and the periods reported in the catalogue. The proportion of targets where the period is within one percent of the reported period is 81.8%. The amount of targets where the period is within a percent of half the reported period is 8.1%, while other close multiplicative factors together constitute 1.3%, adding to a total of 91.2%. Cases where half the period is designated as the orbital period consist of systems with a primary and secondary eclipse that look very similar, and thus create only weak evidence in Fourier space and in phase dispersion for the correct multiple of the period. Figure 12 shows the fractional differences between measured and catalogue periods, around the zero point: we see a very sharp central peak and wide but low tails extending to around $7e-5$ that are imperceptible on the scale of the plot. In the subset with the correct period multiple we see a mean absolute fractional difference of 0.0027% and median of 0.00010%. Similarly, if we divide the difference from the catalogue by the estimated errors instead, as shown in Figure 13, the strong central peak remains but the wings of the distribution become more pronounced, indicating an underestimation of a subset of period errors.

The remaining 8.8% that did not have a multiple of the period consist of three groups, two of which are: very low signal-to-noise cases and those where the binary signal is close to a pure sinusoid, which does not result in a strong series of harmonics to pick up on. The third set of contributors are multiple systems that have more than one period reported in the catalogue: for those systems the strongest periodic signal is found (correctly) and we then miss the other periods as we do not continue to search for higher order multiples.

Eccentricities

The *Kepler* EB Catalog does not provide eccentricity measurements to directly compare against, however, we can present our results in an eccentricity-versus-period plot that gives an idea of whether our measured eccentricity values are globally as we expect them to be. Generally speaking the orbital circularisation time-scale is strongly dependent on the orbital period (Tassoul & Tassoul 1990), which means we would expect all short period binaries to have zero or otherwise small eccentricity. Conversely

⁸ found at keplerebs.villanova.edu




Fig. 12. Histogram and KDE of the fractional difference between the *Kepler* catalogue orbital periods and those measured in this work. Here we select the periods within 1% of the catalogue values. The KDE is scaled to the height of the histogram.

we expect a wide distribution of eccentricities at the high period end as these systems are practically still at their birth configurations.

Figure 14 shows the eccentricity results for a selection of *Kepler* targets. The selection criteria are: an eccentricity measurement error below 0.1 and a catalogue morphology parameter below 0.5, leaving 715 targets of the 2001 systems with an eccentricity measurement. The catalogue morphology parameter is an indicator for the level of detachment between the stars in the system, running from zero to one, with lower values being more detached. We show results that include all morphologies in Figure C.36.

The overall picture is consistent with expectation: we see a buildup of systems at zero eccentricity below around ten days period, and a large range above that period. There are a few clear outliers among the lowest periods that have significant eccentricity measurements, the majority of which is due to an inaccurate identification or measurement of the secondary eclipse. As mentioned in Van Eylen et al. (2016), the measurement of the radial component of eccentricity, $esin(\omega)$, is less straightforward than that of the tangential component, $ecos(\omega)$. This is also apparent for our methodology, which becomes visible when we take a look at the measurements of the tangential component in Figure 15. Apart from the much smaller error estimates, the scatter around zero tangential eccentricity of the cluster below ten days orbital period is tight.

The theory of close binary tides predicts the circularisation time scales of binary orbits, see e.g. Zahn (1975, 1977, 1989); Zahn & Bouchet (1989); Goldreich & Nicholson (1977); Savonije & Papaloizou (1983, 1984). Circularisation is the result of dissipation mechanisms (tidal friction) in the stars interiors, which in turn decrease the orbital energy. Mechanisms include, but are not limited to, damping of the tidal deformation through turbulent flow in a convective region and the damping of tidally




Fig. 13. Histogram and KDE of the orbital period χ -values comparing the *Kepler* catalogue values to those measured in this work. We select the periods within 1% of the catalogue values. The relatively pronounced wings of the distribution indicate underestimated period errors for a portion of systems. The KDE is scaled to the height of the histogram.

forced oscillations through radiative diffusion. The state of circularisation for a single binary depends on its age, which is generally not straightforward to constrain. Instead, the time scales on which binaries of certain types are expected to circularise can be used to form a prediction of what distribution of eccentricities a population of binary systems is expected to cover. Unfortunately, as explained by Van Eylen et al. (2016), age is still hard to completely exclude as a factor for high temperature stars. Due to their much shorter lifespans, comparable to or shorter than the circularisation time scale, they haven't had time to circularise even at short periods where tidal friction is strongest. Additionally, hot stars with radiative envelopes have a different mechanism for circularisation than cool stars with convective envelopes. Therefore we expect hot stars to exhibit larger eccentricities at shorter orbital periods than cool stars, in general.

The *Kepler* sample that Van Eylen et al. (2016) have studied, and that we are looking at here, has a majority of cool stars. If we use the same differentiating temperature of 6250K only one-sixth of the selected eccentricity measurements remain in the high temperature group. In that publication the authors used the tangential component of eccentricity, $ecos(\omega)$ as proxy for eccentricity, whereas we determine both components in our analysis. The two temperature groups look broadly consistent with each other in our measurements, due in part to outliers that may or may not be at correct measurement values. We looked in more detail at some of the clear outliers at low periods with eccentricity measurements at or above 0.1. Most of these measurements arise from misidentifications of noise as a secondary eclipse, or broad ellipsoidal variability as an eclipse (the *Kepler* EB Catalog includes singly eclipsing and even non-eclipsing systems with ellipsoidal variability). Nevertheless, there is a small number of correctly identified systems in this regime, around an ec-




Fig. 14. Eccentricity measurements for a subset of targets from the *Kepler* EB catalogue. Selections were made in eccentricity error (<0.1) and morphology parameter (<0.5), the latter being the more stringent criterion. We use the maximum eccentricity as function of orbital period in Halbwachs et al. (2005, Equation 3) with the cut-off period set to 5 days to plot the dashed line.

centricity of 0.15 with well fitting light curve solutions⁹: specifically KIC numbers 4544587, 7943535, 8196180 and 11867071. Using the temperatures provided by the *Kepler* EB Catalog, all four of these systems fall in the high temperature category, two of them having temperatures above 7000K. This seems to support the conclusion by Van Eylen et al. (2016), who reported an observed difference between the tangential eccentricity distributions of hot-hot binaries as compared to cool-hot and cool-cool binaries, where hot-hot systems are more eccentric at shorter periods. This is in broad lines consistent with theory.

We use a simple formula from Halbwachs et al. (2005, Equation 3) to visualise the envelope of maximum eccentricity in Figure 14 more clearly. The cut-off period, or circularisation period, was set to the value reported by Van Eylen et al. (2016), calculated for cool stars. Apart from the outliers, both those misidentified and the hot stars discussed above, this envelope qualitatively fits the sample. We see the same extension of the low eccentricity cluster towards periods longer than 5 days as in the Van Eylen et al. (2016) sample, which may arise from dependencies on age and/or temperature of the circularisation time scale.

The summarising catalogue of our *Kepler* analysis results is available at the CDS via anonymous ftp to cdsarc.u-strasbg.fr ([insert address]) or via <http://cdsweb.u-strasbg.fr/>.

5. Conclusions

We presented a novel photometric time series analysis methodology and its implementation STAR SHADOW, meant in the first

⁹ And one with a poorer light curve fit that has a high enough tangential component of eccentricity to be sure of its classification in this high eccentricity group.




Fig. 15. Measurements of the tangential component of eccentricity for a subset of targets from the *Kepler* EB catalogue. The subset results from cut-offs in eccentricity error (<0.1) and morphology parameter (<0.5), the latter being the more stringent criterion.

place to analyse large numbers of eclipsing binaries homogeneously, but also broadly applicable for computationally efficient iterative prewhitening. The method is tested against a set of synthetic light curves with known input parameters, and high levels of intrinsic variability other than the eclipses. We find good levels of agreement between the true and extracted sinusoidal frequencies and amplitudes, as well as the true and computed binary parameters. Accuracy and precision varies among the binary parameters, with the tangential component of eccentricity being the most accurate and precise while the radial component is approximately an order of magnitude less well constrained. This is unsurprising as most of the information about the radial component is in the radial velocity curves. The number of extracted frequencies after subtraction of orbital harmonics and after filtering with significance criteria strongly correlates with the number of input sinusoids for the one year long simulated light curves. The number is underestimated for shorter light curves of one month.

We further test the period finding capabilities of our implementation against the *Kepler* EB Catalog, and take a look at the eccentricities obtained for this sample. The method finds the correct multiple or half the orbital period in 90% of cases, showing its capability to be used with minimal predetermined knowledge of the data. Cases where it does find half the period, at less than ten percent, are those where primary and secondary eclipse are nearly identical. Period error estimates are found to be of sufficient size for the vast majority of targets, with heavy tails on the distribution of deviation over error revealing an underestimation for a fraction of targets.

The eccentricities obtained for the *Kepler* targets generally follow the expected pattern as a function of orbital period. Below around ten days eccentricities are broadly consistent with zero and above ten days there is a wide distribution of higher eccentricities. We do see outliers that are caused by misidenti-

cation of the secondary eclipse, due to a combination of its low eclipse depth and the presence of other variability in the out-of-eclipse part of the light curve.

We qualitatively compare our *Kepler* results to the work by Van Eylen et al. (2016), who analysed their own selection of *Kepler* EBs to obtain the tangential part of the eccentricities. We see tentative support for their conclusion that at low orbital periods, hot stars (>6250K) show higher eccentricities than cool stars, which is broadly consistent with the theory. Our eccentricity distribution (for cool stars) looks qualitatively consistent with a cut-off period of 5 days as Van Eylen et al. (2016) compute for cool stars with convective envelopes.

We record an average run time of the analysis of ten minutes for our one year simulated light curves, with an average of two minutes of that spent on the frequency analysis. For the *Kepler* light curves, the recorded average time to completion is just over two hours, and one hour for frequency analysis. The scaling with time series length (four years for *Kepler*) is not linear as the number of extracted sinusoids also increases with time base. This rapid performance is achieved through various optimisations in the algorithm of iterative prewhitening and multi-sinusoid non-linear optimisation, the main ones further described in Appendix B. We encourage others to implement a gradient-based approach for sinusoid optimisation for the free performance boost it gives to save on both time and energy.

Acknowledgements. The research leading to these results has received funding from the KU Leuven Research Council (grant C16/18/005: PARADISE), from the Research Foundation Flanders (FWO) under grant agreements 1124321N (Aspirant Fellowship to LIJ), G089422N (AT), as well as from the BELgian federal Science Policy Office (BELSPO) through PRODEX grant PLATO. [...] This paper includes data collected by the Kepler mission, which are publicly available from the Mikulski Archive for Space Telescopes (MAST) at the Space Telescope Science Institute (STScI). Funding for the Kepler mission is provided by the NASA Science Mission Directorate. STScI is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-26555. Some of the compute resources and services used in this work were provided by the VSC (Flemish Supercomputer Center), funded by the Research Foundation - Flanders (FWO) and the Flemish Government.

Software. This work and the presented code (STAR SHADOW) make use of Python (Python Software Foundation. Python Language Reference, version 3.7. Available at www.python.org) and the Python packages Numpy (Harris et al. 2020), Numba (Lam et al. 2015), Scipy (Virtanen et al. 2020), Astropy (Astropy Collaboration et al. 2013, 2018, 2022), PyMC3 (Salvatier et al. 2016; Salvatier et al. 2016), Theano (The Theano Development Team et al. 2016), Arviz (Kumar et al. 2019), h5py (Collette 2013), Matplotlib (Hunter 2007) and corner (Foreman-Mackey 2016). The authors thank A. Kemp and D. Fritzewski for their valuable input as beta testers of parts of the code presented in this work.

References

Armstrong, D. J., Kirk, J., Lam, K. W. F., et al. 2016, *MNRAS*, 456, 2260
 Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *apj*, 935, 167
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123
 Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33
 Auvergne, M., Bodin, P., Boisnard, L., et al. 2009, *A&A*, 506, 411
 Baran, A. S. & Koen, C. 2021, *Acta Astron.*, 71, 113
 Blanton, M. R., Bershady, M. A., Abolfathi, B., et al. 2017, *AJ*, 154, 28
 Bowman, D. M. & Michielsen, M. 2021, arXiv e-prints, arXiv:2109.10776
 Burkholder, V., Massey, P., & Morrell, N. 1997, *ApJ*, 490, 328
 Collette, A. 2013, Python and HDF5 (O'Reilly)
 De Silva, G. M., Freeman, K. C., Bland-Hawthorn, J., et al. 2015, *MNRAS*, 449, 2604
 Devor, J. 2005, *ApJ*, 628, 411
 Foreman-Mackey, D. 2016, *The Journal of Open Source Software*, 1, 24

Gilmore, G., Randich, S., Asplund, M., et al. 2012, *The Messenger*, 147, 25
 Goldreich, P. & Nicholson, P. D. 1977, *Icarus*, 30, 301
 Guinan, E. F., Ribas, I., Fitzpatrick, E. L., et al. 2000, *ApJ*, 544, 409
 Guo, Z., Ogilvie, G. I., Li, G., Townsend, R. H. D., & Sun, M. 2022, *MNRAS*, 517, 437
 Halbwachs, J. L., Mayor, M., & Udry, S. 2005, *A&A*, 431, 1129
 Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357
 Hocke, K. 1998, *Annales Geophysicae*, 16, 356
 Howard, E. L., Davenport, J. R. A., & Covey, K. R. 2022, *Research Notes of the American Astronomical Society*, 6, 96
 Huber, D. 2015, in *Astrophysics and Space Science Library*, Vol. 408, *Giants of Eclipse: The ζ and Aurigae Stars and Other Binary Systems*, ed. T. B. Ake & E. Griffin, 169
 Hunter, J. D. 2007, *Computing in Science & Engineering*, 9, 90
 IJspeert, L. W., Tkachenko, A., Johnston, C., et al. 2021, *A&A*, 652, A120
 Jennings, Z., Southworth, J., Maxted, P. F. L., & Mancini, L. 2023, *MNRAS*, 521, 3405
 Johnston, C. 2021, *A&A*, 655, A29
 Johnston, C., Aymar, N., Abdul-Masih, M., et al. 2021, *MNRAS*, 503, 1124
 Johnston, C., Pavlovski, K., & Tkachenko, A. 2019a, *A&A*, 628, A25
 Johnston, C., Tkachenko, A., Aerts, C., et al. 2019b, *MNRAS*, 482, 1231
 Johnston, C., Tkachenko, A., Van Reeth, T., et al. 2023, *A&A*, 670, A167
 Kass, R. E. & Raftery, A. E. 1995, *Journal of the American Statistical Association*, 90, 773
 Kirk, B., Conroy, K., Prša, A., et al. 2016, *AJ*, 151, 68
 Koch, D. G., Borucki, W. J., Basri, G., et al. 2010, *ApJ*, 713, L79
 Kollmeier, J. A., Zasowski, G., Rix, H.-W., et al. 2017, arXiv e-prints, arXiv:1711.03234
 Kopal, Z. 1959, *Close binary systems*
 Kumar, R., Carroll, C., Hartikainen, A., & Martin, O. 2019, *The Journal of Open Source Software*, 4, 1143
 Lam, S. K., Pitrou, A., & Seibert, S. 2015, in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, LLVM '15* (New York, NY, USA: Association for Computing Machinery)
 Lampens, P. 2021, *Galaxies*, 9, 28
 Lomb, N. R. 1976, *Ap&SS*, 39, 447
 Lurie, J. C., Vyhmeister, K., Hawley, S. L., et al. 2017, *AJ*, 154, 250
 Massey, P., Morrell, N. L., Neugent, K. F., et al. 2012, *ApJ*, 748, 96
 Maxted, P. F. L. 2016, *A&A*, 591, A111
 Maxted, P. F. L., Gaulme, P., Graczyk, D., et al. 2020, *MNRAS*, 498, 332
 Nelder, J. A. & Mead, R. 1965, *The Computer Journal*, 7, 308
 Pavlovski, K., Southworth, J., Tkachenko, A., Van Reeth, T., & Tamajo, E. 2023, *A&A*, 671, A139
 Press, W. H. & Rybicki, G. B. 1989, *ApJ*, 338, 277
 Prša, A., Batalha, N., Slawson, R. W., et al. 2011, *AJ*, 141, 83
 Prša, A., Guinan, E. F., Devinney, E. J., et al. 2008, *ApJ*, 687, 542
 Prša, A., Kochoska, A., Conroy, K. E., et al. 2022, *ApJS*, 258, 16
 Raftery, A. E. 1995, *Sociological Methodology*, 25, 111
 Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *Journal of Astronomical Telescopes, Instruments, and Systems*, 1, 014003
 Rosu, S., Quintero, E. A., Rauw, G., & Eenens, P. 2023, *MNRAS*, 521, 2988
 Rowan, D. M., Jayasinghe, T., Stanek, K. Z., et al. 2022, *MNRAS*, 517, 2190
 Saha, A. & Vivas, A. K. 2017, *AJ*, 154, 231
 Salvatier, J., Wiecki, T. V., & Fonnesbeck, C. 2016, *PeerJ Computer Science*, 2, e55
 Salvatier, J., Wiecki, T. V., & Fonnesbeck, C. 2016, *ascl:1610.016*
 Savonije, G. J. & Papaloizou, J. C. B. 1983, *MNRAS*, 203, 581
 Savonije, G. J. & Papaloizou, J. C. B. 1984, *MNRAS*, 207, 685
 Scargle, J. D. 1982, *ApJ*, 263, 835
 Sekaran, S., Tkachenko, A., Abdul-Masih, M., et al. 2020, *A&A*, 643, A162
 Sekaran, S., Tkachenko, A., Johnston, C., & Aerts, C. 2021, *A&A*, 648, A91
 Slawson, R. W., Prša, A., Welsh, W. F., et al. 2011, *AJ*, 142, 160
 Southworth, J. 2012, in *Orbital Couples: Pas de Deux in the Solar System and the Milky Way*, ed. F. Arenou & D. Hestroffer, 51–58
 Southworth, J. & Bowman, D. M. 2022, *MNRAS*, 513, 3191
 Southworth, J. & Van Reeth, T. 2022, *MNRAS*, 515, 2755
 Stellingwerf, R. F. 1978, *ApJ*, 224, 953
 Tamuz, O., Mazeh, T., & North, P. 2006, *MNRAS*, 367, 1521
 Tassoul, J.-L. & Tassoul, M. 1990, *ApJ*, 359, 155
 The Theano Development Team, Al-Rfou, R., Alain, G., et al. 2016, arXiv e-prints, arXiv:1605.02688
 Tkachenko, A., Pavlovski, K., Johnston, C., et al. 2020, *A&A*, 637, A60
 Torres, G., Andersen, J., & Giménez, A. 2010, *A&A Rev.*, 18, 67
 Udalski, A., Szymanski, M., Kaluzny, J., Kubiak, M., & Mateo, M. 1992, *Acta Astron.*, 42, 253
 Van Eylen, V., Winn, J. N., & Albrecht, S. 2016, *ApJ*, 824, 15
 Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nature Methods*, 17, 261
 Weidner, C. & Vink, J. S. 2010, *A&A*, 524, A98
 Wells, M. A. & Prša, A. 2021, *ApJS*, 253, 32
 Wilson, R. E. & Devinney, E. J. 1971, *ApJ*, 166, 605
 Wyrzykowski, L., Udalski, A., Kubiak, M., et al. 2003, *Acta Astron.*, 53, 1
 Zahn, J. P. 1975, *A&A*, 41, 329
 Zahn, J. P. 1977, *A&A*, 57, 383
 Zahn, J. P. 1989, *A&A*, 220, 112
 Zahn, J. P. & Bouchet, L. 1989, *A&A*, 223, 112

Appendix A: Additional formulae

Maximum likelihood estimators

Model parameters can sometimes be determined via direct calculation using maximum likelihood estimators (MLEs) under the assumption that our observations are independent and identically distributed according to a normal distribution. We can do this for the two parameters needed to capture linear trends, as well as for higher order polynomials of which we will use the third order here.

We can describe each observation as being drawn from a normal distribution (our first assumption):

$$f(x|\mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(\frac{-(x-\mu)^2}{2\sigma^2}\right), \quad (\text{A.1})$$

where x are the observations, μ is the mean and σ the standard deviation. The likelihood function is the joint probability density of our observations, which under the assumption that our observations are independent, can be written as the multiplication of the individual distributions:

$$L(\theta) = \prod_i f(x_i|\theta) = \left(\frac{1}{\sqrt{2\pi\sigma^2}}\right)^n \exp\left(\frac{-\sum_i(x_i-\mu)^2}{2\sigma^2}\right), \quad (\text{A.2})$$

with θ representing the model parameters and n is the number of data points. An MLE is obtained by maximising the likelihood, or taking the derivative of the likelihood with respect to each model parameter θ and setting this to zero. We take the (natural) logarithm of the likelihood because summation is easier to work with and the location of the maximum is conserved. Another advantage in case of an iterative use of the likelihood is that it is computationally more stable to use the logarithm.

$$l(\theta) \equiv \ln(L(\theta)) = -\frac{n}{2} \ln(2\pi\sigma^2) - \frac{\sum_i(x_i-\mu)^2}{2\sigma^2}. \quad (\text{A.3})$$

We can now confirm that indeed the sample mean ($\frac{1}{n} \sum_i x_i = \bar{x}$) is the MLE of the mean (μ) and the square root of the sample variance ($\sqrt{\frac{1}{n} \sum_i (x_i - \bar{x})^2}$) is the MLE of the standard deviation (σ) by taking the derivative of l to μ and σ respectively and setting this to zero. Substituting these parameters for their estimators, we obtain:

$$l(\theta) = -\frac{n}{2} \ln\left(\frac{2\pi}{n} \sum_i (x_i - \bar{x})^2\right) - \frac{n}{2}. \quad (\text{A.4})$$

The light curve model that we use, whether that is a straight line, sinusoid or more complex, can be seen as the mean of the normal distributions that each observation follows, in the sense that subtracting the model from the data results in a zero mean for all data points. Now let y be our observations and \hat{y} the model, we can then write:

$$l(\theta) = -\frac{n}{2} (\ln(2\pi\sigma_r^2) + 1), \quad (\text{A.5})$$

and σ_r^2 is the variance of the residuals, estimated as:

$$\sigma_r^2 = \frac{1}{n} \sum_i (y_i - \hat{y}_i)^2.$$

We now need an analytical model to work out the estimators for the model parameters. A linear trend is used extensively in this

work, but we also include the quadratic and cubic case here for completeness:

$$\begin{aligned} \hat{y} &= \theta_1 x + \theta_0, \\ \hat{y}_q &= \theta_2 x^2 + \theta_1 x + \theta_0, \\ \hat{y}_c &= \theta_3 x^3 + \theta_2 x^2 + \theta_1 x + \theta_0. \end{aligned}$$

To shorten the following formulae, we define several sums using subscripts denoting the power of the variable within each term as follows:

$$S_{\hat{y}} = \sum_i (y_i - \hat{y}_i)^2 = n\sigma_r^2, \quad (\text{A.6})$$

$$S_{xx} = \sum_i (x_i - \bar{x})^2, \quad (\text{A.7})$$

$$S_{x2x} = \sum_i (x_i^2 - \bar{x}^2)(x_i - \bar{x}), \quad (\text{A.8})$$

$$S_{x2x2} = \sum_i (x_i^2 - \bar{x}^2)^2, \quad (\text{A.9})$$

$$S_{x3x} = \sum_i (x_i^3 - \bar{x}^3)(x_i - \bar{x}), \quad (\text{A.10})$$

$$S_{x3x2} = \sum_i (x_i^3 - \bar{x}^3)(x_i^2 - \bar{x}^2), \quad (\text{A.11})$$

$$S_{x3x3} = \sum_i (x_i^3 - \bar{x}^3)^2, \quad (\text{A.12})$$

$$S_{xy} = \sum_i (x_i - \bar{x})(y_i - \bar{y}), \quad (\text{A.13})$$

$$S_{x2y} = \sum_i (x_i^2 - \bar{x}^2)(y_i - \bar{y}), \quad (\text{A.14})$$

$$S_{x3y} = \sum_i (x_i^3 - \bar{x}^3)(y_i - \bar{y}). \quad (\text{A.15})$$

where \bar{x}^2 is the mean of squares and \bar{x}^3 is the mean of cubes.

Setting $\partial l / \partial \theta_0$ equal to zero for the linear case we find:

$$\theta_0 = \bar{y} - \theta_1 \bar{x}. \quad (\text{A.16})$$

Setting $\partial l / \partial \theta_1$ equal to zero and substituting A.16 we find:

$$\theta_1 = \frac{S_{xy}}{S_{xx}}. \quad (\text{A.17})$$

Maximum likelihood estimators for the standard deviations of the above estimators are:

$$S_{\theta_0}^2 = S_{y,x}^2 \left(\frac{1}{n_{data}} + \frac{\bar{x}^2}{S_{xx}} \right), \quad (\text{A.18})$$

$$S_{\theta_1}^2 = \frac{S_{y,x}^2}{S_{xx}}, \quad (\text{A.19})$$

where the (corrected) standard deviation of $y(x)$ is:

$$S_{y,x}^2 = \frac{S_{\hat{y}}}{n_{dof}}. \quad (\text{A.20})$$

The quadratic case yields:

$$\theta_2 = \frac{S_{x_2y}S_{xx} - S_{xy}S_{x_2x}}{S_{x_2x}S_{xx} - S_{x_2x}^2}, \quad (\text{A.21})$$

$$\theta_1 = \frac{S_{xy} - \theta_2 S_{x_2x}}{S_{xx}}, \quad (\text{A.22})$$

$$\theta_0 = \bar{y} - \theta_2 \bar{x}^2 - \theta_1 \bar{x}. \quad (\text{A.23})$$

For the cubic case, we break up θ_3 in its numerator and denominator:

$$\begin{aligned} \theta_{3,a} = & S_{x_3y} (S_{x_2x}S_{xx} - S_{x_2x}^2) \\ & - S_{x_2y} (S_{x_3x}S_{xx} - S_{x_3x}S_{x_2x}) \\ & + S_{xy} (S_{x_3x}S_{x_2x} - S_{x_3x}S_{x_2x}), \end{aligned}$$

$$\begin{aligned} \theta_{3,b} = & S_{x_3x} (S_{x_2x}S_{xx} - S_{x_2x}^2) \\ & - S_{x_3x} (S_{x_3x}S_{xx} - 2S_{x_3x}S_{x_2x}) \\ & - S_{x_3x}^2 S_{x_2x}, \end{aligned}$$

and obtain the following results:

$$\theta_3 = \frac{\theta_{3,a}}{\theta_{3,b}}, \quad (\text{A.24})$$

$$\theta_2 = \frac{S_{x_2y}S_{xx} - S_{xy}S_{x_2x} - \theta_3 (S_{x_3x}S_{xx} - S_{x_3x}S_{x_2x})}{S_{x_2x}S_{xx} - S_{x_2x}^2}, \quad (\text{A.25})$$

$$\theta_1 = \frac{S_{xy} - \theta_3 S_{x_3x} - \theta_2 S_{x_2x}}{S_{xx}}, \quad (\text{A.26})$$

$$\theta_0 = \bar{y} - \theta_3 \bar{x}^3 - \theta_2 \bar{x}^2 - \theta_1 \bar{x}. \quad (\text{A.27})$$

Bayesian Information Criterion

The BIC of the residuals decreases with more white-noise like residuals (i.e. containing less information), but increases with additional free parameters of which each sinusoid initially has three:

$$BIC = -2 \cdot \ln(L(\theta)) + k \cdot \ln(n), \quad (\text{A.28})$$

where L is the likelihood as function of the parameters θ , n the number of data points and k the number of free parameters. Under the simplifying assumption that the observations are independent and identically distributed according to a normal distribution, the likelihood as in A.5 can be plugged in to get:

$$BIC = n \cdot \ln(2\pi\sigma_r^2) + n + k \cdot \ln(n). \quad (\text{A.29})$$

Kepler's second law

The law of equal swept out areas, or Kepler's second law, couples the orbital properties to time intervals. This means it enables the computation of said properties if we can measure enough 'special' time points along the orbit. With special is meant that we know something about the geometric configuration of the binary at that specific time. For example at the time of eclipse minimum we know that the star in front is covering a maximal portion of the surface area of the star in the back and thus the

projected (in the plane of the sky) distance between their centres is minimal.

For an eccentric orbit, the instantaneous distance between the centres of the two stars is:

$$R = \frac{a(1 - e^2)}{1 + e \cos(\nu)}, \quad (\text{A.30})$$

with a the semi-major axis, e the eccentricity, and ν the true anomaly indicating in what phase of the orbit we are. Plugging this into Kepler's second law in integral form, we obtain:

$$\frac{2\pi(t_b - t_a)}{P} = \int_{\nu_a}^{\nu_b} \frac{R^2 d\nu}{a^2 \sqrt{1 - e^2}} = (1 - e^2)^{\frac{3}{2}} \int_{\nu_a}^{\nu_b} \frac{d\nu}{(1 + e \cos(\nu))^2}, \quad (\text{A.31})$$

where P is the orbital period and $(t_b - t_a)$ is a time interval corresponding to the integral boundary values of ν . As can be read in more detail in Kopal (1959), this integral has an analytic solution:

$$\frac{2\pi(t_b - t_a)}{P} = \left[2 \cdot \arctan \left(\frac{\sqrt{1 - e} \sin\left(\frac{\nu}{2}\right)}{\sqrt{1 + e} \cos\left(\frac{\nu}{2}\right)} \right) - \frac{e \sqrt{1 - e^2} \sin(\nu)}{1 + e \cos(\nu)} \right]_{\nu_a}^{\nu_b}, \quad (\text{A.32})$$

written in a slightly different form than in the aforementioned book. It follows naturally that the time interval is a function of eccentricity and argument of periastron (see Equation 4). What is not immediately apparent from this is the dependence on other properties like inclination and sum of scaled radii. As described in the main text, these properties appear in Equations 5 and 6, which are used to calculate the true anomaly of our special time points that we need to plug into the integral bounds.

An auxiliary angle ϕ_0 is defined corresponding to the sum of scaled radii as in Equation 12, which does not a priori have any physical meaning.¹⁰ This equation replaces the identical factor in Equation 6 to make:

$$\sqrt{1 - \sin^2(i) \cos^2(\phi)} = \sqrt{1 - \sin^2(i) \cos^2(\phi_0)} (1 \pm e \sin(\omega \pm \phi)), \quad (\text{A.33})$$

which is in fact four equations for four different ϕ angles corresponding to first and last contact of the primary and secondary eclipse. Thus, these angles correspond to four additional special time points where we know the orbital configuration. These equations do not have analytical solutions.

In every formula above we assume spherical stars, or rather, circular projections of the stars, which is a decent approximation in the context of eclipses where the stars are always nearly aligned with the line of sight¹¹. It is in the phases in-between the eclipses that we see the stars 'side-on' and their deformation as a result of each others gravity can be significant.

¹⁰ Note that this equation differs a factor $\sqrt{1 - e^2}$ from the one given in Kopal (1959).

¹¹ And stars do not rotate too quickly.

Appendix B: Details of the method implementation

Iterative prewhitening

As already eluded to in the main text, the next frequency, picked by its highest amplitude, is oversampled in frequency space by a factor 100 over the default sampling of the full periodogram. To spend as little computation time as possible on calculating periodograms, we use the ‘fast’ implementation of the Astropy Lomb-Scargle periodogram (Press & Rybicki 1989). This method extrapolates from a fast Fourier transform, and thus is not ideal for high precision. Additionally it cannot be used on sub-domains in frequency space as it will produce unphysical results. However, since we are using a two-step approach, this poses no problems. The oversampling step makes use of the traditional Lomb-Scargle periodogram, albeit an efficient implementation of it by Cuypers, translated by the authors to Python+Numba from the Fortran implementation in the IvS Python Repository (GitHub)¹². The 100 times oversampled periodogram is only computed for the domain between the point to the left and the point to the right of the top most point in the original periodogram (the one at highest amplitude), saving on unnecessary computations. The phase is computed only at the final determined frequency (using the formula by Hocke (1998)), adding very little extra computational cost.

Another optimisation in all of the algorithms where sinusoids are iteratively added to or subtracted from a light curve model is to keep a copy of the model as well instead of just keeping a copy of each sinusoid parameter. Calculating sine waves is very time consuming computationally speaking, so by replacing calculations of sine waves by simple addition and subtraction to the model in memory we save a lot of computation time. Another benefit is that this step of the iterative process no longer scales in time with the number of sinusoids in the model. The additional step of iterating over closely spaced frequencies does still add a level of dependence on the number of sinusoids, making the overall scaling of iterative prewhitening with the number of sinusoids slightly more than linear (with the dependence being a function of frequency density in the periodogram). The overall increase in efficiency is large enough that for the data used in testing for the speed (a fairly short light curve) the process even becomes memory bound instead of CPU bound, meaning that the speed of the computer memory is now the bottleneck for this part of the process in that scenario.

Minimisation with Scipy

Nelder-Mead (NM) (Nelder & Mead 1965) is a robust gradient free method that was tested extensively in this framework, and it performs well in terms of final objective function value reached. However, it requires many function evaluations and thus is computationally expensive. The optimisation steps in the algorithm are (were) the computational bottleneck, so we have looked for an effective way of decreasing this computational cost. Scipy offers a wide range of minimisation methods: in search of a faster but still robust method in our application we tested all of them. Some methods require a Jacobian or both Jacobian and Hessian to be provided, and several cannot handle boundaries. For our application of fitting a sum of sine waves to the light curve, we can compute the partial derivatives constituting the Jacobian and Hessian, and our testing described below refers to this application. Table B.1 below summarises the features of each minimisation method.

¹² Our translation is tested to be identical in speed.

Table B.1. Scipy minimize method features.

Method name	Bounds	Constraints	Requires Jacobian	Requires Hessian	Can use Jacobian	Can use Hessian
Nelder-Mead	■					
Powell	■					
CG				■		
BFGS				■		
Newton-CG		■		■	■	
L-BFGS-B	■			■		
TNC	■			■		
COBYLA		■				
SLSQP	■	■		■		
trust-constr	■	■		■	■	
dogleg			■	■	■	■
trust-ncg			■	■	■	■
trust-exact	■		■	■	■	■
trust-krylov	■		■	■	■	■

Notes. The methods of the function `scipy.optimize.minimize` have different features and requirements in terms of using the Jacobian or Hessian matrix and in being able to handle bounds or constraints.

Although it is not strictly necessary for every minimisation step in the algorithm, we do want to impose boundaries on our parameters. We tested methods that have the option to numerically compute gradients with and without the analytical Jacobian to see what the performance gain would be. Both Powell and SLSQP (without Jacobian) are more than an order of magnitude faster than NM, but both diverge from the true solution instead of converging on it. SLSQP becomes slightly faster when the Jacobian is provided and diverges slightly less dramatically. Method trust-constr (without Jacobian) takes a factor of a few longer than NM, while reaching worse cost function values. However, when provided with the Jacobian it redeems itself by becoming faster than NM by a factor two and obtaining the same cost function value. Unfortunately providing the Hessian as well makes it much slower again; this could be due to the high computational cost of the Hessian function. Methods CG, BFGS, L-BFGS-B and TNC perform similar to NM in terms of cost function, but only BFGS and L-BFGS-B are faster by a noticeable factor of a few without the Jacobian. CG does not benefit time-wise from the Jacobian, but the three others do: they reach similar speeds of more than an order of magnitude faster than NM. Also Newton-CG, which required the Jacobian to function, falls in the same performance category as those three mentioned in both cost function and speed. It does not benefit meaningfully from the Hessian. All four methods that require the Hessian take a longer time than NM did by a factor more than two. Since CG, BFGS and Newton-CG cannot handle bounds nor constraints, L-BFGS-B and TNC go on to a second round of testing. Also COBYLA makes it to the second round by being fast and only marginally worse in terms of cost function, and its lack of support for bounds can be mitigated by its support for constraints.

In a more demanding test with a longer light curve, the COBYLA method (which does not use a Jacobian) shows comparable speed to L-BFGS-B and TNC (with Jacobian), but diverges and does not return a sensible result. L-BFGS-B and TNC (both compared with and without Jacobian) perform very similarly: the final choice for L-BFGS-B is made based on its faster

speed, and even though the cost function value it reaches is not quite as good, the difference in the actual light curve model is imperceptibly small.

Analytical gradients

As mentioned above, the optimisation makes good use of analytical gradient information. We include their mathematical form here. Our objective function is the negative log-likelihood:

$$f = -\ln(l) = \frac{n}{2} \left[\ln \left(\frac{2\pi}{n} \sum_i (y_i - \hat{y}_i)^2 \right) + 1 \right], \quad (\text{B.1})$$

with symbol definitions as in Appendix A. The Jacobian is defined as:

$$J = \left[\frac{\partial f}{\partial \theta_1}, \frac{\partial f}{\partial \theta_2}, \dots, \frac{\partial f}{\partial \theta_n} \right], \quad (\text{B.2})$$

where each θ_j stands for one of the model parameters. Using the generalised product rule the derivatives can be written as:

$$\frac{\partial f}{\partial \theta_j} = \sum_i \frac{\partial f}{\partial \hat{y}_i} \frac{\partial \hat{y}_i}{\partial \theta_j}, \quad (\text{B.3})$$

$$\frac{\partial f}{\partial \hat{y}_i} = \frac{-n(y_i - \hat{y}_i)}{\sum_i (y_i - \hat{y}_i)^2}, \quad (\text{B.4})$$

$$\frac{\partial \hat{y}_i}{\partial \theta_j} = \frac{\partial}{\partial \theta_j} a_j \sin(2\pi f_j(t_i - t_{mean}) + \phi_j). \quad (\text{B.5})$$

Making the derivatives explicit, we get:

$$\frac{\partial \hat{y}_i}{\partial f_j} = 2\pi(t_i - t_{mean}) a_j \cos(2\pi f_j(t_i - t_{mean}) + \phi_j), \quad (\text{B.6})$$

$$\frac{\partial \hat{y}_i}{\partial a_j} = \sin(2\pi f_j(t_i - t_{mean}) + \phi_j), \quad (\text{B.7})$$

$$\frac{\partial \hat{y}_i}{\partial \phi_j} = a_j \cos(2\pi f_j(t_i - t_{mean}) + \phi_j), \quad (\text{B.8})$$

and for the orbital period:

$$\frac{\partial \hat{y}_i}{\partial p_{orb}} = \sum_h 2\pi \frac{-n_h}{p_{orb}^2} (t_i - t_{mean}) a_h \cos \left(2\pi \frac{n_h}{p_{orb}} (t_i - t_{mean}) + \phi_h \right), \quad (\text{B.9})$$

in which we are now summing over all harmonic sinusoids.

Note that for the eclipse model parameters we cannot compute gradients analytically. However, this does not prevent us from using the gradients that we do know during the optimisation of the eclipse model plus sinusoids. We use the `scipy.optimize.approx_fprime` to obtain numerical gradients for the 6 parameters of the eclipse model and fill those in at their respective locations in the Jacobian. This provides a similar, although smaller, speed increase as for the pure sinusoid optimisations.

Parameter space

The parameter space we work in for the six free parameters in our eclipse model is made up of: $ecos(\omega)$, $esin(\omega)$, $cos(i)$, ϕ_0 , $\log_{10} \left(\frac{r_2}{r_1} \right)$ and $\log_{10} \left(\frac{sb_2}{sb_1} \right)$. This choice of parameters ensures two things: minimal correlation between the free parameters and a flatter space across the relevant domain. In the case of $ecos(\omega)$, $esin(\omega)$ and ϕ_0 both of these reasons apply, while the reduction of correlation is the driving factor compared to their counterparts e , ω and r_{sum} . Reducing correlation has a strong effect on the ability of optimisation algorithms to accurately and quickly converge. Another benefit is that our correlation structures, and thus error estimates, in the least correlated parameter space are most reliable and simple to interpret.

Parameters $\frac{r_2}{r_1}$ and $\frac{sb_2}{sb_1}$ have an extremely curved parameter space in the sense that a step of 0.1 has a completely different meaning depending on where in the parameter domain we are situated. Using logarithms, $\log_{10} \left(\frac{r_2}{r_1} \right)$ and $\log_{10} \left(\frac{sb_2}{sb_1} \right)$, has two main advantages that are similar to the ones mentioned for the reduction in correlation. It makes it simpler for many optimisation algorithms to traverse the parameter space for a more accurate and sometimes more speedy result, as well as improving reliability and interpretability of the error estimates (in that parameter space).

Translating eclipse timings to eccentricity

Here, we describe the steps for obtaining physical parameters from time measurements in more detail.

Although a multitude of approaches was tested, we settled on one that provides the best combination of speed and accuracy. The initial step is to set $cos(i) = 0$ and compute the approximate ϕ_0 with Equation 11. Sufficient accuracy is retained when solving Equation 7 and plugging in the approximate Equations 8 and 9. Iteratively updating ϕ_0 using exact formulae for the eclipse durations does however provide a noticeable improvement in the measurement of the sum of scaled radii through Equation 12.

To find inclination, ratio of radii and ratio of surface brightness we construct a cost function based on iteratively solving Equations 5 and A.33 and plugging the obtained angles into 4 and then A.32 to obtain accurate theoretical eclipse times, and computing theoretical depths with Equation 13. Minimisation is done with a global optimiser from Scipy, called `simplicial`¹³ homology global optimisation (SHGO), with reasonable boundaries on the parameters for feasible systems.

Minimisation is performed in the parameter space: $cos(i)$, $\log_{10} \left(\frac{r_2}{r_1} \right)$, $\log_{10} \left(\frac{sb_2}{sb_1} \right)$. This has the benefit of homogenising the step size across the parameter space, where especially the two parameter ratios suffer from a large difference in step size on either end of their physical ranges. Since we also compute components of eccentricity $ecos(\omega)$ and $esin(\omega)$, we also need to convert them to e and ω :

$$e = \sqrt{(e \cos(\omega))^2 + (e \sin(\omega))^2}, \quad (\text{B.10})$$

$$\omega = \arctan \left(\frac{e \sin(\omega)}{e \cos(\omega)} \right),$$

where code-wise an `arctan2` function is used instead of `arctan` to obtain correct angles across the full range. ϕ_0 is converted to the sum of radii with:

$$r_{sum} = \sqrt{1 - \sin^2(i) \cos^2(\phi_0)} (1 - e^2). \quad (\text{B.11})$$

¹³ Indeed, `simplicial`, not `simplicitic`.

Error estimation is done with a combination of analytical error formulae and the Monte Carlo approach of importance sampling. The analytical formulae are used as minimum error estimates, but we do not have approximate analytical formulae for all parameters. Furthermore, for the analytical error formulae, we do not have an a priori value for the inclination error, so we estimate one that is informed by the tests on synthetic light curves. Starting with an error estimate of 0.035 radians in inclination (2 degrees), we use the following formulae as errors on the given subset of parameters:

$$\begin{aligned}\sigma_{e\cos(w)} &= \sqrt{\partial_p^2 \sigma_p^2 + \partial_i^2 (\sigma_{t,1}^2 + \sigma_{t,2}^2) + \partial_i^2 \sigma_i^2}, \\ \partial_p &= \frac{-\pi(t_2 - t_1)\sin^2(i)}{p_{orb}^2(1 + \sin^2(i))}, \\ \partial_i &= \frac{\pi\sin^2(i)}{p_{orb}(1 + \sin^2(i))}, \\ \partial_i &= \frac{2\sin(i)\cos(i)}{(1 + \sin^2(i))^2},\end{aligned}\quad (B.12)$$

where we abbreviated the partial derivatives: $\partial_x \equiv \frac{\partial_x}{\partial_x}$.

$$\begin{aligned}\sigma_{e\sin(w)} &= \sqrt{\partial_p^2 \sigma_p^2 + \sum_{m,n} \partial_\tau^2 \sigma_{\tau,m,n}^2 + \partial_i^2 \sigma_i^2 + \partial_{\phi,0}^2 \sigma_{\phi,0}^2}, \\ \partial_p &= \frac{-\pi(\tau_{1,2} + \tau_{1,1} - \tau_{2,1} - \tau_{2,1})\sin^2(i)\cos^2(\phi_0)}{2\sin(\phi_0)p_{orb}^2(1 + \sin^2(i)(1 + \cos^2(\phi_0)))}, \\ \partial_\tau &= \frac{\pi\sin^2(i)\cos^2(\phi_0)}{2\sin(\phi_0)p_{orb}(1 - \sin^2(i)(1 + \cos^2(\phi_0)))}, \\ \partial_i &= \frac{2\sin(i)\cos(i)\cos^2(\phi_0)}{(1 - \sin^2(i)(1 + \cos^2(\phi_0)))^2}, \\ \partial_{\phi,0} &= \frac{2\sin^2(i)(1 - \sin^2(i))\sin(\phi_0)\cos(\phi_0)}{(1 - \sin^2(i)(1 + \cos^2(\phi_0)))^2},\end{aligned}\quad (B.13)$$

where the sum goes over primary ($m = 1$) and secondary ($m = 2$) ingress ($n = 1$) and egress ($n = 2$).

$$\sigma_e = \sqrt{\cos^2(w)\sigma_{e\cos(w)}^2 + \sin^2(w)\sigma_{e\sin(w)}^2},\quad (B.14)$$

$$\sigma_w = \sqrt{\frac{\sin^2(w)}{e^2}\sigma_{e\cos(w)}^2 + \frac{\cos^2(w)}{e^2}\sigma_{e\sin(w)}^2},\quad (B.15)$$

$$\begin{aligned}\sigma_{\phi,0} &= \sqrt{\partial_p^2 \sigma_p^2 + \sum_{m,n} \partial_\tau^2 \sigma_{\tau,m,n}^2}, \\ \partial_p &= \frac{\pi(\tau_{1,2} + \tau_{1,1} + \tau_{2,1} + \tau_{2,1})}{2p_{orb}^2}, \\ \partial_\tau &= \frac{\pi}{2p_{orb}},\end{aligned}\quad (B.16)$$

$$\begin{aligned}\sigma_{r,sum} &= \sqrt{\partial_e^2 \sigma_e^2 + \partial_i^2 \sigma_i^2 + \partial_{\phi,0}^2 \sigma_{\phi,0}^2}, \\ \partial_e &= 2e\sqrt{1 - \sin^2(i)\cos^2(\phi_0)}, \\ \partial_i &= \frac{\sin(i)\cos(i)\cos^2(\phi_0)(1 - e^2)}{\sqrt{1 - \sin^2(i)\cos^2(\phi_0)}}, \\ \partial_{\phi,0} &= \frac{\sin^2(i)\sin(\phi_0)\cos(\phi_0)(1 - e^2)}{\sqrt{1 - \sin^2(i)\cos^2(\phi_0)}},\end{aligned}\quad (B.17)$$

The error formulae for $e\sin(w)$ and ϕ_0 are found lacking in an earlier version of our synthetic test (Section 3), which likely is because they both rely on the measurement of the widths of the eclipses, which can be inaccurate and is more easily disturbed by additional variability. Based on the same test, we construct simple scaling relations that apply to the minimum error in the concerned parameters. Their functional form is:

$$\text{minimum error} = C \cdot \exp\left(-\frac{\text{depth}_2}{0.3 \sum_h a_h}\right),\quad (B.18)$$

where the constant C is 0.2 for $e\sin(w)$ and 0.05 for ϕ_0 . The specific scaling variable used in the exponent was found to correlate strongest with absolute deviation from truth of these parameters (correlation of -0.45 for $e\sin(w)$), much more so than secondary eclipse depth alone (correlation of -0.24 for $e\sin(w)$). The inclusion of the depth of the secondary eclipse is intuitive as shallower secondary eclipses are harder to measure. The other parameter is less intuitive: it is the sum over amplitudes of low and high frequency harmonics. The sum actually goes over harmonic orders up to and including 5, and harmonic orders above 15. Dependence on the higher frequency harmonics may be understood in terms of the presence of additional variability in the light curve, which reduces the accuracy of the measurements of eclipse edge points. The dependence on the low frequency harmonics is likely to do with a stronger presence of these harmonics indicating either very wide eclipses or out-of-eclipse variability, which cross over into the regime where the simplified physics of the applied formulae starts to lose accuracy.

To obtain minimum error values for the two ratios and their logarithms, we compute a scale factor by dividing the computed value for σ_e by the error estimate obtained for e in the importance sampling (explained below). We take the maximum of the upper and lower error value for this, and multiply the importance sampling errors of the logarithm of radius ratio and the logarithm of surface brightness ratio, again taking the maximum of upper and lower error value. Finally, computing the relative error scaling in log space, and multiplying by the respective parameters in linear space, we obtain the minimal errors for both representations of these two parameters.

In the importance sampling, we generate input normal distributions of a thousand samples for the measured values of eclipse times (minima, contact, tangency) within their respective error regions and compute the output of the iterative ‘translation’ scheme outlined above for each of the generated input vectors. This results in a set of output distributions for each computed parameter. Estimates of the errors are then obtained using the highest density interval (HDI) at 0.683 probability. Note that the resulting error region can exclude the point estimate made before. This is fixed by imposing the minimum errors. The choice not to use the mean of the distribution as point estimate is made because the distributions are often multi-modal and may not accurately represent the optimal solution.

Appendix C: Additional plots

Synthetic test set

These additional Figures aim to give a more complete view of the output of our analysis of the synthetic light curves discussed in Section 3. The first, Figure C.1, shows the same period measurements as Figure 4, except now ordered by primary eclipse depth. We note that the eclipse depth does not have any significant influence on the period accuracy except in the most extreme cases where no eclipses can be made out (shaded and hatched dark grey area).

The majority of the remaining figures show the absolute and relative distributions of the parameter outputs compared to the true values. The relative distributions (denoted with χ) use the estimated error values obtained from the importance sampling and after imposing our error minima (from the formal error calculations and scaling relations). The two colours, blue and orange, in each plot show values obtained at the eclipse timing translation step and the eclipse model fitting step, respectively.

Finally we show obtained inclinations versus the third light that was added to the artificial light curves. Third light is not modelled, but Figure C.27 shows no sign of a correlation between third light and inclination. These parameters are mutually degenerate, so such a correlation would be expected: possibly the overall (low) accuracy of our inclination measurements hides this dependence.




Fig. C.1. Blue points with error bars: fractional difference between the measured and the true orbital period with simple linear regression error estimates, sorted by descending primary eclipse depth. Grey line with points: primary eclipse depths, plotted on the right axis. Note that cases shaded grey and hatched do not show eclipses, and cases shaded orange have fewer than three eclipse cycles.




Fig. C.2. Histogram and KDE of the eccentricity deviations from the truth.

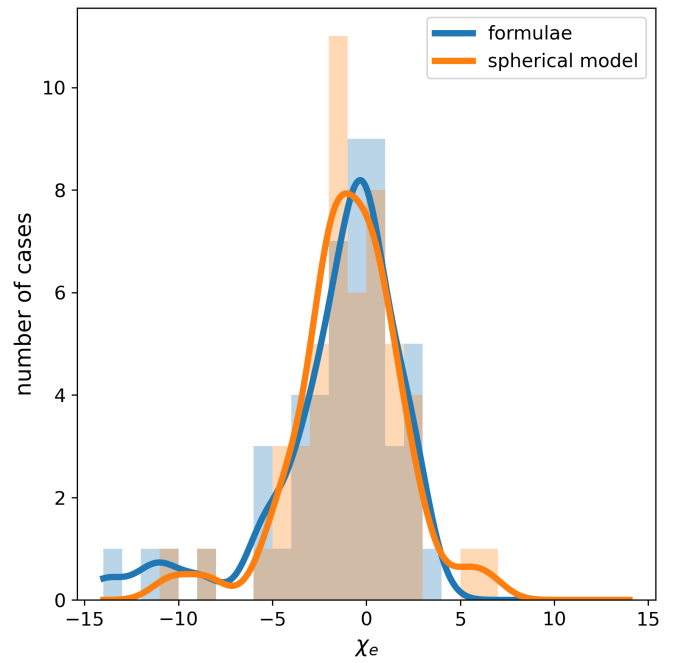


Fig. C.3. Histogram and KDE of the eccentricity χ -values (deviation from the truth divided by error estimate).




Fig. C.4. Histogram and KDE of the argument of periastron deviations from the truth.




Fig. C.6. Histogram and KDE of the inclination deviations from the truth.




Fig. C.5. Histogram and KDE of the argument of periastron χ -values (deviation from the truth divided by error estimate).




Fig. C.7. Histogram and KDE of the inclination χ -values (deviation from the truth divided by error estimate).




Fig. C.8. Histogram and KDE of the sum of scaled radii deviations from the truth.




Fig. C.10. Histogram and KDE of the ratio of radii deviations from the truth.




Fig. C.9. Histogram and KDE of the sum of scaled radii χ -values (deviation from the truth divided by error estimate).




Fig. C.11. Histogram and KDE of the ratio of radii χ -values (deviation from the truth divided by error estimate).




Fig. C.12. Histogram and KDE of the ratio of surface brightness deviations from the truth.




Fig. C.14. Histogram and KDE of the tangential component of eccentricity deviations from the truth.




Fig. C.13. Histogram and KDE of the ratio of surface brightness χ -values (deviation from the truth divided by error estimate).




Fig. C.15. Histogram and KDE of the tangential component of eccentricity χ -values (deviation from the truth divided by error estimate).




Fig. C.16. Histogram and KDE of the radial component of eccentricity deviations from the truth.




Fig. C.18. Histogram and KDE of the cosine of inclination deviations from the truth.




Fig. C.17. Histogram and KDE of the radial component of eccentricity χ -values (deviation from the truth divided by error estimate).




Fig. C.19. Histogram and KDE of the cosine of inclination χ -values (deviation from the truth divided by error estimate).




Fig. C.20. Histogram and KDE of the auxiliary angle ϕ_0 deviations from the truth.




Fig. C.22. Histogram and KDE of the logarithm of ratio of radii deviations from the truth.




Fig. C.21. Histogram and KDE of the auxiliary angle ϕ_0 χ -values (deviation from the truth divided by error estimate).




Fig. C.23. Histogram and KDE of the logarithm of ratio of radii χ -values (deviation from the truth divided by error estimate).




Fig. C.24. Histogram and KDE of the logarithm of ratio of surface brightness deviations from the truth.

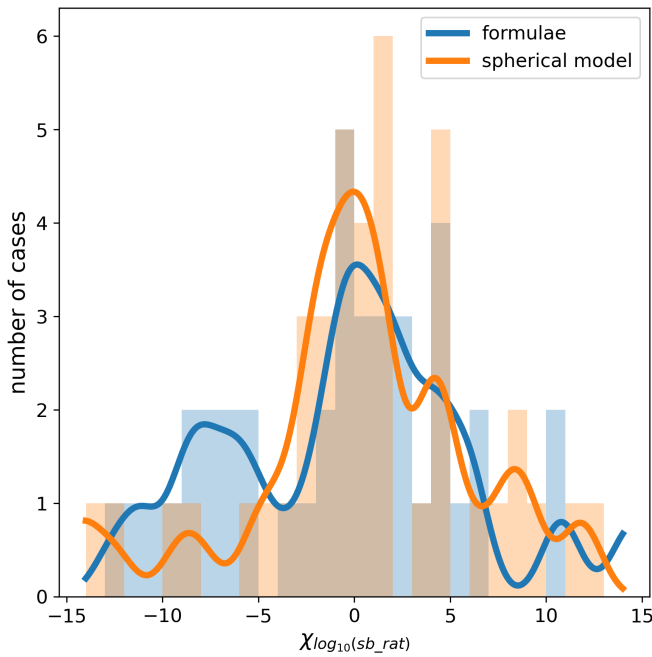


Fig. C.25. Histogram and KDE of the logarithm of ratio of surface brightness χ -values (deviation from the truth divided by error estimate).




Fig. C.26. Tangential part of eccentricity (top panel) and deviation from truth (bottom panel) versus secondary eclipse depth.




Fig. C.27. Inclinations (top panel) and deviation from truth (bottom panel) versus third light (fraction of total light). Note the absence of an apparent correlation between accuracy of inclination and third light.

Example light curves

Here we show a few examples of the analysis results from the synthetic set of test light curves discussed in Section 3. The average level of added sinusoid amplitudes is high in this set, challenging the algorithm for eclipse detection and measurement. In Figure C.28 the algorithm was successful in detecting and measuring the eclipses, perhaps with the exception of a remaining residual spike on the egress of the primary.

We see better how the added variability affects the measurement in Figure C.32, where the primary is wider to one side than it should be. The resulting large residuals are captured by the sinusoids that are complementary to the eclipse model to make up the full variability of the light curve. This results in harmonic sinusoids that in this case model unphysical/artificial signal, but in other cases can also help model for example the physical shape of the eclipses or ellipsoidal variability not captured by the eclipse model (see Figure C.33).




Fig. C.28. Example (full) synthetic light curve of case 7. The top panel shows the final eclipse model of the eclipses in red and the full model of the light curve including sinusoids in grey. The bottom panel shows the residuals of subtracting the eclipse model (blue) and those of subtracting the full model (orange). The model of sinusoids is also plotted separately in red here. The slight overestimation of the primary eclipse width can be identified by sharp residual peaks here and leads to an underestimation of the radial part of eccentricity of about 0.2 for this case.




Fig. C.29. The Fourier amplitude spectrum of case 7. In red and pink are the extracted sinusoids, of which the red ones qualify as harmonics. At an orbital period of 16.56, the spacing of harmonics is quite dense (grey grid lines).




Fig. C.30. Example synthetic light curve of case 99 folded over the orbital period. The top panel shows the final eclipse model of the eclipses in red and the full model of the light curve including sinusoids in grey. The orange model is the starting point for the light curve fit. The bottom panel shows the residuals of subtracting the eclipse model (blue) and those of subtracting the full model (orange). The model of sinusoids is also plotted separately in red here. The model starts from an overestimated primary eclipse width, and the fit is unable to fully correct for this as the residuals show. This leads to an underestimated tangential eccentricity by about 0.2 for this case.




Fig. C.31. The Fourier amplitude spectrum of case 7 after subtracting the eclipse model. In red are the extracted sinusoids that pass all criteria, including the signal to noise threshold shown in green (calculated in a window of $1d^{-1}$).




Fig. C.32. Example (full) synthetic light curve of case 23. The top panel shows the final eclipse model in red and the full model of the light curve including sinusoids in grey. The bottom panel shows the residuals of subtracting the eclipse model (blue) and those of subtracting the full model (orange). The model of sinusoids is also plotted separately in red here. We can see that the primary eclipse is too wide as the detection algorithm picked up on other variability close to it.




Fig. C.33. Example (folded over the period) synthetic light curve of case 24. The top panel shows the final eclipse model in red and the full model of the light curve including sinusoids in grey. The bottom panel shows the residuals of subtracting the eclipse model (blue) and those of subtracting the full model (orange). The model of sinusoids is also plotted separately in red here. The eclipse model has trouble fitting these wide eclipses and does not capture any out-of-eclipse variability.




Fig. C.34. Example (folded over the period) synthetic light curve of case 88. The top panel shows the final eclipse model in red and the full model of the light curve including sinusoids in grey. The bottom panel shows the residuals of subtracting the eclipse model (blue) and those of subtracting the full model (orange). The model of sinusoids is also plotted separately in red here.




Fig. C.35. Histogram and KDE of the orbital period χ -values comparing orbital periods measured in this work to those from the *Kepler* catalogue. Here we select the periods within 1% of half the catalogue values and the errors are taken to be the SLR uncertainties. The KDE is scaled to the height of the histogram.




Fig. C.36. Eccentricity measurements for a subset of targets from the *Kepler* EB catalogue. A cut-off was made in eccentricity error (<0.1) but all morphology values were included. We use the maximum eccentricity as function of orbital period in Halbwachs et al. (2005, Equation 3) with the cut-off period set to 5 days to plot the dashed line.

Kepler EB Catalog

Figure C.35 is complementary to Figure 13: it shows the distribution of measured periods around half the catalogue values. The shape of the distribution differs little from the one around the catalogue period, with a sharp central peak and heavy tails.

Adding all morphologies from the *Kepler* EB catalogue to the eccentricity-period plot, Figure C.36, extends it substantially to lower periods. We now include 1399 of 2001 targets with an eccentricity measurement, as compared to the 715 we had when limiting to morphologies below 0.5. We see that the majority of additional systems fall at around zero eccentricity, and while the error bars are relatively large, there is more scatter than for the previous selection of targets below ten days orbital period. Especially noticeable are the systems at period around a day and high eccentricity. These do not indicate physical peculiarities, but rather inaccurate determination of eclipse times. For these low period systems, eclipses are often not sharply defined making it harder for our derivative-based method to accurately determine the edges.

We see from Figure C.37 that most of these outlier systems are not due to large $e \cos(\omega)$ values: even though the scatter around the zero line for the tangential component of eccentricity is markedly larger for these high morphology targets, it stays within a value of 0.1 for most of them. That means inaccurate $e \sin(\omega)$ measurement is the driving factor for these outliers.




Fig. C.37. Measurements of the tangential component of eccentricity for a subset of targets from the *Kepler* EB catalogue. The subset results from a cut-off in eccentricity error (<0.1) but all morphology values were included.

Appendix D: Correlation reduction

Reducing the correlation between optimisation parameters can strongly benefit the ease and accuracy with which the optimisation is done. It may be possible to transform a parameter to this end; one such example is the use of auxiliary angle ϕ_0 as mentioned in the main text. For the reason that we are trying to solve the general case of an eclipsing binary light curve, it may not always make sense to use a certain reparametrisation. For example: when eclipses are grazing, instead of showing the flattened off phase of either total eclipses or transits, we lose information about the eclipsing stars. In this case a grazing eclipse provides less information on the radii, but we do not lose all. It may then be appealing to keep the ratio of radii fixed at one, while only fitting for sum of scaled radii. We can, however, do better. The only information we actually lose completely is which star is bigger and which is smaller, but we can still recover information about the size of the bigger star and the size of the smaller star. We can represent this loss of knowledge in a new free parameter, which replaces the ratio of radii:

$$\rho = \frac{r_1^2 + r_2^2}{a^2}, \quad (\text{D.1})$$

where r_i is each star's radius and a the semi-major axis, similar to the sum of radii. We keep the sum of radii:

$$r_{sum} = \frac{r_1 + r_2}{a}, \quad (\text{D.2})$$

or alternatively angle ϕ_0 may be chosen. We cannot recover the individual radii r_1 and r_2 from this combination of parameters, representing the lack of information in the light curve. What we can recover is:

$$r_{s,l} = \frac{1}{2}r_{sum} \pm \frac{1}{2}\sqrt{2\rho - r_{sum}^2}, \quad (\text{D.3})$$

with which we find a small and a large radius (implicitly scaled by a). Clearly for the systems where we have the information available, we would not want to artificially limit ourselves, in addition to actually increasing correlations between the parameters in those cases. Therefore we choose to accept sub-optimal correlation structures in return for a more generally and homogeneously applicable method. However, for analysis into a specific binary with grazing eclipses (which is a large fraction of all eclipsing binaries for geometrical reasons), this approach may aide the specific optimisation procedure utilised, increasing accuracy and potentially speed.