

A new methodology for classifying eclipsing binaries with *Kepler* data and deep learning

Mousam Mondal¹★, Patricia Cruz², Hugh R. A. Jones¹, M. C. Gálvez-Ortiz² and John F. Aguilar^{3,4}

¹Centre for Astrophysics Research, University of Hertfordshire, College Lane, Hatfield, Hertfordshire AL10 9AB, UK

²Centro de Astrobiología (CAB), CSIC-INTA, Camino Bajo del Castillo s/n, E-28692 Villanueva de la Cañada, Madrid, Spain

³Departamento de Matemáticas, Universidad Militar Nueva Granada, kilometro 2 vía Cajicá–Zipaquirá, código postal 110111, Colombia

⁴PhD Programme in Astrophysics, Doctoral School, Universidad Autónoma de Madrid, Ciudad Universitaria de Cantoblanco, E-28049 Madrid, Spain

Accepted 2026 June 16. Received 2026 June 9; in original form 2025 October 21

ABSTRACT

We present a new method for the automated classification of eclipsing binaries, into contact, detached, and semidetached types using *Kepler* data. Phase-folded light curves are generated and chi-square versus box size plots are constructed by comparing flux values to the median flux, revealing distinct class patterns. These patterns were first modelled using a polynomial damped sinusoidal function, whose period served as classification feature, achieving an overall accuracy of 86.5 per cent. To capture more features and enhance accuracy, we trained a convolutional neural network, which improved the total accuracy to 90 per cent, including 47 per cent for the challenging semidetached systems. However, several binaries displayed irregular chi-square signatures. To mitigate this, we incorporated simulated light curves generated with the PHOEBE modelling code, achieving 99 per cent accuracy in distinguishing contact and detached binaries. The resulting chi-square morphologies show a strong correlation with orbital period, and a subset of systems exhibit quarterly variability in their light curves and chi-square trends. We designate these as temporally varying systems. By measuring the normalized spread of the chi-square period across quarters, we define a statistical threshold that separates these systems from stable binaries. We reported four temporally varying systems not previously noted in the literature with magnetic activity that requires further investigation. Furthermore, cooler stars, namely late-F, G, K, and M types, display systematically higher variability than hotter stars. Cross-matching with catalogues of magnetically active stars indicates that stellar flares and starspots are the most likely causes of this enhanced variability.

Key words: eclipsing binaries – methods: data analysis – methods: statistical – techniques: photometric – stars: variables: general – Machine Learning.

1 INTRODUCTION

The advancement of space technology and modern telescopes has led to an abundance of high-precision astronomical data. Visual inspection remains a reliable approach for eclipsing binaries (EB) classification, capturing subtle morphological features in photometric light curves (LCs) that are difficult to encode algorithmically. However, its manual nature makes it impractical due to the rapidly growing volume of survey data. This created increasing opportunities for automated methods and artificial intelligence algorithms to identify patterns, anomalies, and correlations in a fast and efficient way. Among all the astronomical targets, EBs offer great potential for pattern recognition and classification based on their LCs. The analysis of LCs is crucial for the derivation of physical parameters, such as the relative radii of their components, as well as their orbital properties, which can serve as test-beds to improve atmospheric models, formation, and evolution theories (e.g. J. Andersen 1991; G. Torres et al. 2010; W. Frost et al. 2024).

EBs can be classified according to the shape of their LCs or their physical properties. Based on physical configuration, three main morphological EB categories are defined by the degree to which the component stars fill their Roche lobes (J. Kallrath & E. F. Milone 2009). Detached systems have both components within their respective Roche lobes. In semidetached systems, only one component fills this boundary. Contact and overcontact binaries occur when both components exceed this limit and share a common envelope.

In the past years, many researchers have employed automated methods to classify EBs using machine learning models trained and validated on large data sets of LCs. For instance, L. M. Sarro et al. (2006) developed an automated classification approach using

★ E-mail: mousamphysics137@gmail.com

a Bayesian ensemble of neural networks trained on data from the *Hipparcos* satellite. Their method classified both pulsating and EBs based solely on LC morphology and further subdivided EBs into four geometrically meaningful classes.

Using a collection of long-term time-series, G. Matijević et al. (2012) classified EBs in the second *Kepler* catalogue (R. W. Slawson et al. 2011) using a non-linear dimensionality reduction tool known as Locally Linear Embedding (LLE) (e.g. S. T. Roweis & L. K. Saul 2000; S. Ridder & R. Duin 2002). Using this method, LC data in higher dimensions were projected into lower 2D space. A classification parameter, c , was included to provide the fit between 2D LC and spline function. They did not report a direct classification accuracy, but instead introduced the c parameter, whose ranges can be mapped to detached, semidetached, and contact categories.

M. Čokina et al. (2021) used deep learning techniques for EB classification. Their model classified semidetached EBs as detached systems and achieved an accuracy of 98 per cent, as it was primarily trained on detached and overcontact EBs. When semidetached EBs were excluded, the accuracy improved to 100 per cent. More recently, I. V. Daza-Perilla et al. (2023) employed a supervised machine learning model with a compound decision tree (CDT) to classify EBs into the same classification categories, utilizing VISTA (Visible and Infrared Survey Telescope for Astronomy) in the Vía Láctea (VVV; D. Minniti et al. 2010; R. K. Saito et al. 2012; M. Hempel et al. 2014) and the VISTA Variables in the Vía Láctea eXtended Survey (VVVX; D. Minniti et al. 2018) survey. The study identified 35 features, such as shape and period. The CDT model reliability was approximately 80 per cent when compared to visually performed classifications. In their recent work, X. Ding et al. (2025) developed a neural-network-accelerated PHOEBE framework combined with Markov chain Monte Carlo (MCMC) to model semidetached eclipsing binaries from *TESS* light curves, identifying 327 systems in which the less massive component fills its Roche lobe and three systems in which the more massive component does so. The neural network replaced PHOEBE within the MCMC loop, providing nearly two orders of magnitude faster light-curve generation.

In this work, we propose a different approach, an indirect classification method, based on the LC normalization technique outlined in P. Cruz et al. (2022), to classify systems into contact EBs (CEBs), semidetached EBs (SDEBs), and detached EBs (DEBs). For simplicity, we grouped CEBs and overcontact systems together under CEBs, maintaining a classification structure with three primary categories. This paper is organized as follows. The process of sample selection is described in Section 2, while the LC normalization procedure is explained in Section 3. Section 4 provides detailed information about the classification scheme, including the use of two different methodologies. Section 5 contains the results and discussion, the summary and our conclusions are presented in Section 6.

2 DESCRIPTION OF THE SAMPLE

To develop our automated classification scheme, a data set of well-sampled LCs with thousands of epochs each was needed. Therefore, we adopted the complete sample of 2920 EBs from the *Kepler* Eclipsing Binary Catalogue¹ (KEBC; B. Kirk et al. 2016). In the *Kepler* mission, observations are divided into seventeen observing quarters, each corresponding to an approximately 93 d observing segment defined by 90° spacecraft rolls required to maintain spacecraft orientation (R. L. Morris et al. 2020). To collect accurate LC data for each quarter for all targets, we utilized the Mikulski Archive for Space Telescopes (MAST)².

From this sample, we executed a query on the MAST *Kepler* Search portal, setting the condition flag to EB as the target type, and searched for long cadence data within a 0.02 arcmin radius. We specifically selected targets with long-cadence (1765.5 s) LCs to maximize our sample size, as short-cadence (58.85 s) data were available for only a limited number of targets. Subsequently, we downloaded quarter-wise fits files for all selected targets to facilitate further analysis.

After cross-matching the KEBC with photometric data from the MAST, we selected 2865 EB systems with long-exposure data suitable for our automated classification scheme. The orbital periods (P_{orb}) of these targets, taken from the KEBC, ranged from 0.073 to 1087.3 d (Fig. 1). We then examined the P_{orb} distribution for each visually classified EB class to determine the period range of CEBs, SDEBs, and DEBs. The accuracy of our automated classification depends on the visual inspection of the 2865 LCs, where we identified EBs into three morphological categories: CEBs, SDEBs, and DEBs, as described in detail in Section 4. Based on this cross-match, we found that CEBs had P_{orb} ranging from 0.073 to 7.056 d, DEBs from 0.184 to 1087.3 d, and SDEBs from 0.216 to 5.499 d. In this work, CEBs, overcontact systems and ellipsoidal variables were treated as a single CEB category, preserving a three-class structure. We did not treat ellipsoidal variables as a separate class in this work; therefore, any systems showing ellipsoidal-type variability, including the few long-period objects classified as CEBs, are retained within the CEB category.

3 NORMALIZATION PROCEDURE

The automated classification scheme was developed as a by-product of the LC normalization procedure adopted by P. Cruz et al. (2022). In this work, we generated normalized phase-folded LCs for all objects in our sample using the same methodology, as only around 1000 objects were analysed in P. Cruz et al. (2022). For consistency, we also adopted the P_{orb} values given by the KEBC for all 2865 objects in our sample. To prepare the data, we first conducted a median-based normalization using the pre-search data conditioning (PDC) flux, where the PDC flux values were divided by the respective median for each quarter; the PDC flux was produced by the *Kepler* pipeline to

¹<https://keplerEBs.villanova.edu>

²<https://archive.stsci.edu>




Figure 1. Box plots illustrated the P_{orb} values of C, SD, and D EBs. Each box represented the interquartile range (IQR), which encompassed 50 per cent of the data. The horizontal line within each box indicated the median P_{orb} . The whiskers extended to 1.5 times the IQR, and the dots above the whiskers represented outliers – data points that fell outside this range, highlighting extreme values in each category. In the figure, CEBs had 98 outliers (11.88 per cent), SDEBs had 12 outliers, and DEBs had 269 outliers (14.04 per cent). The error bars, represented by the whiskers, indicated the spread of the data within 1.5 times the IQR.

remove common instrumental trends, such as pointing drifts and thermal effects, present in the raw simple aperture photometry (SAP) flux, while preserving the dominant astrophysical variability (S. J. Murphy 2012; J. E. Van Cleve et al. 2016).

Normalization was carried out quarter-by-quarter using the sliding-median method of P. Cruz et al. (2022). For each quarter, we constructed a local median baseline by moving a window across the LC. In the classification plots, we used the half-window length as the box size and explored box sizes from 10 to 100 epochs, corresponding to total sliding windows of 21 to 201 epochs. For each trial window size, an error-weighted chi-square value was calculated by comparing the observed flux with the corresponding sliding-median baseline, following the procedure described by P. Cruz et al. (2022). The optimal window size was then selected using the same criterion: the chi-square values were scanned in order of increasing window size, and the first value at which the integer part of the chi-square stopped increasing monotonically was adopted, where the integer part was used to suppress noise-level fluctuations that would otherwise cause a premature break in the scan. This selected the smallest window that could follow the dominant long-term trend while avoiding excessive smoothing of the signal.

The sliding-median baseline obtained with the selected window size was then smoothed using a cubic spline. The original LC was normalized by dividing it by this spline-smoothed trend, and the normalized quarterly LCs were then combined to produce the final normalized LC for each target.

The chi-square–box-size relation showed how the LC responded to filtering on progressively longer time-scales. The *Kepler* long-cadence sampling interval used in this work was 1765.5 s, corresponding to approximately 29.43 min. In terms of the plotted box-size scale, a box size of 10 epochs corresponded to about 4.9 h, while a box size of 100 epochs corresponded to about 49.0 h, or roughly 2 d.

The chi-square–box-size curve was used in two ways in this work. First, it helped us choose the normalization window for each quarter by identifying the smallest box size that could follow the dominant instrumental trend. Secondly, by visually inspecting the chi-square plots together with the phase-folded LCs, we found that systems in the same morphological class generally showed similar chi-square trends. DEBs usually produced an increasing, polynomial-like or exponential-like curve (Fig. 2, panels a–c), whereas CEBs showed a damped sinusoidal pattern (Fig. 2, panels g–i). The SDEB systems showed less distinct behaviour; most displayed a distorted sinusoidal or damped pattern between the CEB and DEB cases (Fig. 2, panels d–f). These class-dependent differences made the chi-square–box-size curves useful diagnostic features for separating the three EB morphological classes.

For a given target, the chi-square curves were generally similar across different observing quarters, indicating that this diagnostic was stable over the *Kepler* baseline. Since the number of data points varied from quarter to quarter, we used the quarter with the largest number of epochs for the subsequent classification analysis.

To test whether the selected range was robust, we also calculated chi-square–box-size curves up to 1000 epochs for a representative subset of systems. For CEBs, the damped sinusoidal structure remained visible up to about 180–200 epochs. Beyond this range, the curves became almost flat. This flattening occurred because larger windows covered multiple orbital cycles, so the sliding median averaged over all orbital phases and approached a nearly constant baseline. DEBs showed a simpler behaviour. Their chi-square values increased at small box sizes, approximately below 25 epochs, and then remained nearly constant for larger windows. For long-period DEBs, this early plateau is expected to be related to the eclipse duration rather than the full P_{orb} , because individual eclipses are narrow compared with P_{orb} . Once the window became wider than the characteristic eclipse width, the sliding-median response saturated. This was also seen in our tests up to 1000 epochs, corresponding to approximately 20.4 d for the *Kepler* long-cadence sampling used here, where the DEB chi-




Figure 2. The chi-square versus box size plots for nine different EBs are shown. The top row (a, b, c panels) displays three different DEBs with *Kepler* IDs (KIC): 1026032, 1873513, and 2166200, respectively. The middle row (d, e, f panels) displays three different SDEBs with KIC: 2577756, 5560831, and 5962514, respectively. Lastly, the bottom row (g, h, i panels) displayed three CEBs systems with KIC: 1433410, 2012362, and 2159783, respectively.

square behaviour showed no significant change beyond approximately 25 epochs. Since the main aim of this work was morphological classification rather than optimal detrending of individual systems, we restricted the main analysis to box sizes of 10–100 epochs. This range captured the relevant class-dependent structure while remaining applicable to the full sample.

To verify the consistency of the above-described behaviour, we analysed the original LCs, using PDC flux values available in the *Kepler* FITS files, as well as the normalized LCs. The analysis was performed for both, individual quarter data and the complete LCs which included all quarters. In most cases, we noticed similar trends, also suggesting a chi-square plot dependence on the P_{orb} of the EBs. The major discrepancies were observed in the chi-square plots of some SDEBs and short P_{orb} DEBs, even for different quarters of observations for a given object, which we discussed further in Section 5.

4 MORPHOLOGICAL CLASSIFICATION

Our classification techniques comprised two processes: the curve fitting approach, described in Section 4.1, and the deep learning approach, described in Section 4.2.

During LCs visual inspection, we identified EBs whose LCs did not fall under the three standard morphological classes (C, SD, D). These were systems that did not show any primary or secondary eclipses visible in the phase-folded LC with the given P_{orb} (e.g. KIC 9512958). Earlier authors R. W. Slawson et al. (2011) and G. Matijević et al. (2012) mentioned these systems in the *Kepler* EB catalogues and termed them as uncertain categories (UNC). R. W. Slawson et al. (2011) included UNC as a morphological category. Also, the KEBC (B. Kirk et al. 2016) catalogue included various types of EB LC, which do not satisfy our three classes: heartbeat stars (e.g. KIC 5944240), UNC classes (e.g. KIC 5823121), pulsation EBs (e.g. KIC 7515679), and eclipse timing variations. We classified all these ambiguous cases under the UNC category and they were excluded from our analysis.




Figure 3. Random Forest feature-importance ranking of the six free parameters of the PDS function for distinguishing CEBs, DEBs, and SDEBs. The vertical axis lists the six fit parameters: ω – the angular frequency; C – the constant vertical offset parameter of the fit; ϕ – the phase angle of the sinusoidal component; B – the exponential damping coefficient; A – the overall amplitude of the function; and n – the exponent of the polynomial term. The horizontal axis gives the normalized feature importance of each parameter, while the colour scale represents the corresponding relative importance, with brighter colours indicating a higher contribution to the Random Forest classifier.

After visually classifying the LCs, we found that out of 2865 objects, there were 832 CEBs, 1820 DEBs, 154 SDEBs, and 59 UNC systems. This brought our total visually classified sample size to 2806, which contained only the three morphological classes. This classification was used for comparison with the automated classification result.

4.1 Classification using the PDS function

Our first approach was based on analysing the shape of the chi-square plots and finding the best-fitting function. Since polynomial or damping-like trends shown in the plots of DEBs and CEBs dominated our chi-square plots (95 per cent of the sample), with SD systems showing a similar trend or something in between, we adopted the following polynomially damped sinusoidal (PDS) function.

$$y = A \cdot x^n \cdot \exp(-Bx) \cdot \sin(\omega x + \phi) + C, \quad (1)$$

where A is the amplitude of the function, ω is the angular frequency, and ϕ is the phase angle. The B reflected the decay rate in the exponential decline of the oscillatory component with respect to the box size x in our chi-squared plots. The parameter C is a constant vertical offset parameter that shifts the fitted curve along the y-axis and accounts for any residual baseline level in the normalized chi-square values. Additionally, x^n polynomial modulation term, where n is the polynomial exponent.

We utilized Python's `Scipy curve fit` (e.g. P. Virtanen et al. 2020) code for optimizing parameters and fitting functions to data. We tried fitting the chi-square values directly without the offsets, scaling, and alternative non-linear optimization strategies; however, these approaches did not yield accurate recovery of the oscillatory PDS component across the full sample and were unable to reproduce the sinusoidal behaviour. To simplify the fitting procedure, we normalized the data, which scaled the data values between $[0, 1]$ and then the PDS function was fitted to the normalized data, with the baseline offset absorbed by the constant offset parameter in equation (1), which accounts for any residual vertical shift during fitting. As outlined in equation (1), the PDS function encompassed multiple parameters crucial for accurately capturing the system's behaviour.

To examine which PDS parameters most effectively distinguished the three EB classes, we trained a Random Forest classifier (L. Breiman 2001) using the 2806 visually classified systems in our sample. The six fitted PDS parameters were used as input features. The classifier achieved a fivefold cross-validation accuracy of 87.6 ± 2.8 per cent, and the corresponding feature-importance ranking was shown in Fig. 3. Among the parameters, ω emerged as the strongest discriminator with an importance of 0.340, while the constant vertical offset parameter ranked second at 0.238. These normalized scores summed to unity and represented each parameter's fractional share of the classifier's total discriminative power across the three EB classes, with a score of 1 indicating a feature that did all the work alone and 0 indicating a feature that contributed nothing. Constant offset parameter was mainly associated with separating D systems from the other two classes. This was expected, since D systems often produced more polynomial-like chi-square plots, whereas C and SDEBs showed more continuously varying morphologies and therefore overlapped more strongly in baseline level. Since the broad separation between D and the other two systems was already largely captured by P_{PDS} , we used the coefficient of determination (R^2) as a complementary diagnostic of the PDS fit quality. In this context, R^2 acted as a shape indicator: high values corresponded to the well-behaved chi-square plots typically seen in C and D systems, while lower values highlighted the distorted plots more commonly associated with SD systems.

The PDS period P_{PDS} , used as our principal physical descriptor throughout this work, was obtained directly from the fitted ω as $P_{\text{PDS}} = 2\pi/\omega$. Since ω determined the spacing of extrema in box size, this period estimate was independent of the amplitude, baseline, or normalization applied to the chi-square plots.




Figure 4. Chi-square plot examples of three EB classes: (a) DEB (KIC 8299947); (b) SDEB (KIC 2577756); and (c) CEB (KIC 1433410). The chi-square values are Min–Max normalized to the range [0, 1]. Data points are shown as filled circles, and the dashed curve shows the PDS function fitting.

Table 1. Different P_{PDS} values related to visually classified C and D systems and their ratios, along with the average ratio and ratio difference.

P_{PDS} value	C Below	D Above	C Ratio (per cent)	D Ratio (per cent)	Average ratio ($\bar{R}$, per cent)	Ratio difference (per cent)
10	434	1816	52.16	99.78	75.97	47.62
14	560	1809	67.31	99.40	83.35	32.09
18	620	1793	74.52	98.52	86.52	24.00
22	663	1778	79.69	97.69	88.69	18.00
26	674	1759	81.01	96.65	88.83	15.64
30	699	1734	84.01	95.27	89.64	11.26
34	714	1696	85.82	93.19	89.50	7.37
38	726	1653	87.26	90.82	89.04	3.56
42	733	1624	88.10	89.23	88.67	1.13
46	747	1598	89.78	87.80	88.79	1.98
50	755	1548	90.75	85.05	87.90	5.69

After selecting the initial fitting values and identifying P_{PDS} as the primary classification feature, our code proceeded with the fitting using the curve-fit method. We then adopted R^2 as an additional fit-quality metric to evaluate how well the PDS function described each chi-square curve. We used R^2 as a scale-independent indicator of how well the PDS function captured the shape of the normalized chi-square versus box-size plots; prior to fitting, the chi-square curves were rescaled to the range [0,1] to enable consistent comparison across targets. The R^2 scores for both DEBs and CEBs were consistently greater than 0.9, indicating good fits, as shown in Figs 4(a) and (c), respectively. On the other hand, SDEBs showed comparatively poorer fits than the other two classes, as illustrated in Fig. 4(b), where most SDEBs had R^2 values below 0.9, often by a substantial margin. Accordingly, while P_{PDS} effectively separated CEBs and DEBs, R^2 was adopted as a complementary diagnostic to identify SD systems, since the PDS model provided a comparatively poor description of their chi-square plots. The combined use of P_{PDS} and R^2 therefore enabled a more consistent classification across all three EB classes.

Initially, we conducted our analysis on a small sample of 100 EBs. We found that CEBs generally exhibited shorter P_{PDS} values compared to DEBs. To select the best P_{PDS} value for classification, we first calculated how many systems visually classified as DEBs or CEBs fell under certain threshold values of P_{PDS} . In Table 1, we showed different P_{PDS} values and, corresponding to each value, the percentage of visually identified CEBs that fell below the threshold, the percentage of visually identified DEBs that exceeded the threshold, as well as their average ratio percentage and ratio difference percentage. Our selection criteria focused solely on maximizing overall accuracy without favouring either class. Because DEBs and CEBs accounted for approximately 95 per cent of our sample, we based the criteria on the classification performance of these two classes, iterating the defined thresholds as follows. In Fig. 5(a), we presented the classification accuracy of CEBs (in black) and DEBs (in blue) for different P_{PDS} values using an initial combined sample of 100 systems. The accuracy was based on visual versus automated comparisons, with the automated classifications using different P_{PDS} thresholds shown in Fig. 5(a) and the R^2 criteria listed in Table 2. The green line in this plot represented the total average accuracy of CEBs and DEBs, and the red line represented the difference between C and D accuracy. We wanted to choose the optimum P_{PDS} value for classification by comparing P_{PDS} values from 10 to 50. We wanted to select a threshold P_{PDS} value that was not biased toward any class. For example, if we selected a threshold P_{PDS} value of 18, it would have given a CEB accuracy of 74 per cent but a DEB accuracy of 98.5 per cent. However, as DEBs dominated the sample, the overall accuracy would have increased, but we aimed to select a P_{PDS} value that represented both classes equivalently with the highest accuracy. This is why the P_{PDS} value of 30 was selected as the optimal choice. At this threshold, the overall misclassification rate, defined as $100 - \bar{R}$, where $\bar{R}$ represents the average classification ratio of CEBs and DEBs, reached its lowest value of 10.36 per cent at $P_{\text{PDS}} = 30$, corresponding to an average classification ratio of 89.64 per cent. This




Figure 5. (a) Plot shows how different P_{PDS} values can separate visually identified CEBs (black) and DEBs (blue). The green dashed line represented the average classification accuracy for both CEBs and DEBs, the red dashed line indicated the accuracy difference between CEBs and DEBs, and the purple vertical line marked the threshold value of 30, which was the best P_{PDS} value for classification. (b) The x-axis in the second plot represented CEBs (black) and DEBs (blue) under two distinct P_{PDS} thresholds: ≤ 10 and ≥ 50 . The y-axis illustrated the percentage of CEBs and DEBs in the entire sample, highlighting how their distribution varied across these thresholds. This comparison emphasized how each P_{PDS} range captured specific EB types effectively.

Table 2. Adopted criteria in our PDS-fitting method.

Classification	Fitting criteria
Contact	$(R^2 \geq 0.9 \text{ and } P_{\text{PDS}} < 30) \text{ or } P_{\text{PDS}} \leq 10$
Detached	$(R^2 \geq 0.9 \text{ and } P_{\text{PDS}} \geq 30) \text{ or } P_{\text{PDS}} \geq 50$
Semidetached	$(R^2 < 0.9)$

confirmed that $P_{\text{PDS}} = 30$ provided the best balance between the two dominant classes. Furthermore, the CEBs accuracy reached 84.01 per cent, while DEBs accuracy remained high at 95.27 per cent. This combination provided a reasonable balance, ensuring that neither class was disproportionately prioritized, as shown in Table 1. The P_{PDS} value of 30 thus represented the best compromise between high overall accuracy and fairness in class representation. The purple vertical line depicted P_{PDS} value 30 in Fig. 5(a). Along with this, in Fig. 5(b), we also showed that when P_{PDS} value was 10 or less, EBs were primarily CEBs, with only 0.15 per cent of samples being DEBs. When P_{PDS} value was 50 or higher, EBs were primarily DEBs, with only 2.9 per cent of samples being CEBs.

Finally, we used a combination of both R^2 and P_{PDS} to distinguish between CEB and DEB systems. The final classification criteria were presented in Table 2. Compared with the visual inspection, out of 832 CEBs, 680 were correctly classified by the PDS function, yielding an accuracy of 81.7 per cent. Out of 1820 DEBs, 1714 were classified as D, yielding an accuracy of 94.1 per cent. Finally, out of 154 SDEBs, only 34 were classified as such, yielding a poor accuracy of 22.1 per cent. Had we prioritized SDEB accuracy, our overall performance would have suffered significantly; therefore, we chose a method that was biased toward the two dominant classes (CEBs and DEBs) and incorporated the R^2 criterion to recover some SDEBs. Considering the whole sample, the total accuracy obtained with the PDS-function fitting was 86.5 per cent.

Table 3 presented a portion of our classification results, showing first 10 systems from each morphological class (C, SD, D) in ascending order based on their *Kepler* ID. The complete version of this table, listing all 2865 targets, is available online.

Analysing the misclassified cases, we found some EBs with clearly different class assignments when comparing the visual classification of the phase-folded LC and the chi-squared plot fitting procedure, as well as targets with curves that could not be clearly discriminated. Since the PDS function was an attempt to explain the chi-squared plot behaviour and, in general, it fitted the SDEBs poorly, the results were not ideal and presented lower accuracy than initially expected. Furthermore, the P_{PDS} values are directly correlated with the binary P_{orb} , which ultimately represents a separation of periods rather than a purely morphological classification feature, as discussed further in Section 5. Therefore, we adopted a deep-learning technique to address the classification problem.

4.2 Classification using a deep-learning method

We employed the convolutional neural network (CNN) approach to perform more in-depth pattern recognition on chi-square plots and to enhance the previous classification accuracy across three classes. We utilized a two-step approach. First, we normalized the chi-square data to bring it to the same scale, which helped in analysing the meaningful structure of the plots rather than their varying magnitudes. For this purpose, we utilized a Min–Max normalization technique, which scaled the data values between $[0, 1]$. This normalization was particularly suitable for the adopted 1D-CNN model, as neural networks generally performed better when input data was standardized to a fixed range, ensuring faster convergence and improved training stability (F. Pedregosa et al. 2011).

Before selecting the final architecture, we evaluated multiple pattern-recognition approaches, including Fourier-based features, wavelet-transformed representations, and higher dimensional CNN variants (e.g. 2D and 3D CNNs). While these methods captured

Table 3. The table presents the classification of the first 10 CEBs, DEBs, and SDEBs that are common across all three methods and are sorted by their *Kepler* ID in ascending order. The *Kepler* ID column represents the unique Kepler identification number, while the Visual column indicates the visual classification of those targets. The R^2 column denotes the goodness-of-fit of the PDS function, and the P_{PDS} column lists the periods derived from the PDS method. The PDS Classification column refers to the classification output from the PDS fitting method. The prediction of CNN1 column shows the classification probability from a 1D Convolutional Neural Network trained on *Kepler* all-quarter light-curve data, and CNN1 represents the corresponding classification output. Similarly, prediction of CNN2 provides the classification probability from a 1D-CNN model trained on both *Kepler* data and synthetic data generated using PHOEBE, with the final classification output shown in the CNN2 column. The full version of this table is available online.

<i>Kepler</i> ID	Visual classification	R^2	P_{PDS}	PDS Classification	Prediction of CNN1	CNN1	Prediction of CNN2	CNN2
KIC 002437038	C	0.988	6.489	C	0.957	C	0.915	C
KIC 002571439	C	0.987	7.787	C	0.969	C	0.879	C
KIC 002715417	C	0.996	5.725	C	0.907	C	0.924	C
KIC 002717141	C	0.998	9.560	C	0.889	C	0.758	C
KIC 002858322	C	0.816	10.643	S.D	0.526	S.D	0.643	C
KIC 003127873	C	0.996	16.302	C	0.737	C	0.529	C
KIC 003221207	C	0.997	11.489	C	0.717	C	0.943	C
KIC 003232823	C	0.992	10.264	C	0.943	C	0.773	C
KIC 003353233	C	0.988	7.664	C	0.969	C	0.840	C
KIC 003437800	C	0.995	8.796	C	0.918	C	0.878	C
KIC 001026957	D	0.998	211.866	D	0.956	D	0.792	D
KIC 001571511	D	0.993	41.430	D	0.946	D	0.872	D
KIC 002438490	D	0.998	122.812	D	0.962	D	0.796	D
KIC 002576692	D	0.995	56.430	D	0.931	D	0.631	D
KIC 002719873	D	0.952	78.341	D	0.945	D	0.941	D
KIC 002860788	D	0.894	1960.258	D	0.964	D	0.934	D
KIC 002987433	D	0.961	13695137.660	D	0.946	D	0.852	D
KIC 003102024	D	0.960	140.488	D	0.930	D	0.922	D
KIC 003122985	D	0.995	77.939	D	0.730	D	0.640	D
KIC 003230227	D	0.993	16355.969	D	0.869	D	0.896	D
KIC 002305372	S.D	0.906	32.274	D	0.887	D	0.932	D
KIC 002577756	S.D	0.743	20.478	S.D	0.670	S.D	0.425	C
KIC 003730067	S.D	0.991	7.114	C	0.961	C	0.868	C
KIC 005077994	S.D	0.990	16.814	C	0.719	S.D	0.417	C
KIC 005962514	S.D	0.988	37.378	D	0.590	C	0.631	D
KIC 006046061	S.D	0.116	25941846.164	D	0.969	C	0.832	C
KIC 006220497	S.D	0.987	31.425	D	0.717	S.D	0.912	D
KIC 006387887	S.D	0.995	10.510	C	0.929	C	0.776	C
KIC 006443392	S.D	0.931	18.629	C	0.776	S.D	0.539	C
KIC 006516874	S.D	0.505	29450204.538	D	0.860	S.D	0.549	S.D

certain global characteristics of the sequences, they did not provide consistent improvements in validation accuracy and often increased model complexity or training instability.

Given that the Chi-square curves represent structured 1D sequences, a 1D CNN provides a natural inductive bias by learning local morphological patterns directly from the ordered data. 1D-CNNs have been widely and successfully applied in astronomical time-series analysis, including light-curve classification (M. Čokina et al. 2021) and solar flare prediction (P. Li et al. 2025), where they effectively capture local temporal features while maintaining computational efficiency. Motivated by these successful 1D-CNN applications, we adopted a compact 1D-CNN with two Conv1D layers per branch and a lightweight three-branch ensemble, which provided stable convergence and strong performance without resorting to deeper architectures.

The model followed a three-branch 1D-CNN design implemented through two custom functions. The first function (`model_architect`) built a single-scale 1D-CNN with five convolutional layers (kernel size of 8, tanh activation), followed by a global max-pooling layer that captured the most informative local patterns across the input sequence. The pooled output was passed through a dense layer with 50 units and then a dropout layer with a rate of 10 per cent to reduce overfitting. The second function (`core_model`) instantiated three parallel branches with identical architecture, each processing the same input sequence. The resulting feature representations were concatenated into a 150D vector and passed to a dense softmax classifier to predict the probabilities of the three classes (C, SD, D). Training used the Adam optimizer (D. P. Kingma & J. Ba 2014) together with a focal loss function $\gamma = 2$, $\alpha = 0.25$ (T.-Y. Lin et al. 2020), which reduced the impact of easily classified examples and focused on harder ones. To prevent overfitting, early stopping with a patience of 20 epochs was applied. The model was trained with a batch size of 64, which provided a balance between efficiency and stable gradient updates, and a maximum of 100 epochs to ensure convergence without excessive training. For clarity and reproducibility, the complete CNN architecture, including all layers and hyperparameters, is summarized in Table 4. The above complete CNN algorithm is available on GitHub³.

³<https://github.com/mousam100>

Table 4. Summary of the CNN architecture and training configuration used for morphological classification of EBs.

Architecture component	Architecture details
Input representation	Each sample is a chi-square-boxsize time-series stored as a sequence of length 91 with two channels (box size and chi-square value), giving an input shape of (91,2).
model_architect (single branch)	Returns a Keras sub-model that maps an input sequence of shape (91,2) to a 50D embedding: Conv1D (50 filters, kernel size = 8, tanh, same padding) → Conv1D (×4, 50 filters, kernel size = 8, tanh, same padding) → GlobalMaxPooling1D → Dense (50, tanh) → Dropout (0.1).
core_model (multibranch assembly)	Instantiates three parallel branches by calling model_architect three times and applying each branch to an input of shape (91,2). The three branch outputs are concatenated and passed to a softmax classifier.
Branch configuration	Three branches with identical layer structure and hyperparameters. In the implemented code, branches are instantiated separately (i.e. each branch has its own learnable weights).
Merge	Concatenation of three embeddings ($3 \times 50 = 150$).
Classifier	Dense (3, softmax).
Loss function	Focal loss ($\gamma = 2, \alpha = 0.25$).
Optimizer	Adam.
Training setup	Batch size = 64; maximum epochs = 100; early stopping (patience = 20, monitored on validation loss, best weights restored).

4.2.1 CNN with Kepler data

CNN required a large amount of data to train effectively. However, in our *Kepler* sample, we only had 2806 EB systems. Therefore, we used the chi-square plots of all the quarters from all systems to increase our data set. Chi-square plots for most targets exhibited similar patterns across quarters. Here, we used a labelled train-test split, where the quarters with the highest number of epochs, the same analysed in Section 4.1, were included in the testing set only and were randomly selected. For training, we used chi-square plots from the remaining sample; we have also seen that some chi-square plots show contradiction to their majority shape, and we have taken out these targets from the training set and only selected those matching Fig. 2 and which were consistent across all quarters.

To maintain the randomness of the split and avoid any biased results, a Python-based random sampling algorithm was applied. After excluding 9 per cent of the EB sample whose chi-square plots showed morphologies inconsistent with their assigned classes, the final data set consisted of 48 per cent of the objects for training and 43 per cent for testing. These EBs were excluded to help the model learn representative patterns of the majority classes, and we discussed these systems in more detail in Section 5. The test set contained 1220 EB systems, including 800 DEBs, 350 CEBs, and 70 SDEBs, and the 48 per cent training set contained 21 564 chi-square plots in total: 14 370 plots from 900 DEBs, 6639 plots from 423 CEBs, and 555 plots from 36 SDEBs. From this training set, we further selected 20 per cent (4313 plots) for validation during model training. Since the chi-square plots were mixed across systems, it was not possible to specify how many individual EBs contributed to the validation subset, but the proportion of chi-square plots remained 20 per cent of the training set.

Our model achieved an overall accuracy of 90.3 per cent when classifying into three categories. Among the 800 DEBs, 750 were correctly classified, giving an accuracy of 93.8 per cent. For the 350 CEBs, 319 were accurately identified, resulting in an accuracy of 91.1 per cent. Lastly, from the 70 SDEBs, 33 were correctly detected, yielding an accuracy of 47 per cent. These results were depicted in the confusion matrix in Fig. 6(a).

Some of the SD chi-square plots had significant similarities with those from C and D systems. Therefore, in our second experiment, we classified only CEBs and DEBs, excluding the SDEBs, to assess the model's accuracy without the uncertainty introduced by SD systems. We used the same labelled train – test and validation sample proportion method as in the previous experiment and kept the same DEB and CEB sample sizes, excluding only the SDEBs. A total of 14 370 chi-square plots from 900 DEBs and 6639 plots from 423 CEBs were used for training. In the test data, out of 350 CEBs, 335 were correctly classified, which resulted in an accuracy of 95.7 per cent. Similarly, out of 800 DEBs, 753 were correctly classified, which gave an accuracy of 94.1 per cent. This brought the overall classification accuracy to 95 per cent. The confusion matrix for this classification was shown in Fig. 6(b).

The precision, recall, and F1-score metrics offer valuable insights into the accuracy of a CNN model's classification of the test sample. Precision measures the accuracy of true positive predictions out of all positive predictions, with a high precision score indicating a high likelihood of true positive prediction. Recall, also known as the true positive rate, gauges the precision or true positive rate out of all positive instances. The F1 score balances precision and recall, serving as the harmonic mean of these two metrics. Together, these three scores provide information about the model's performance and its effectiveness in classification tasks. In Table 5, we presented the Precision, Recall, and F1-scores for these two analyses. The columns labelled C (CNN1), D (CNN1), and SD (CNN1) represent the metrics using only the *Kepler* data, where the first three rows of CNN1 correspond to the first analysis – considering all three morphological classes – and the last three rows correspond to the analysis with only CEBs and DEBs.

One of the possible reasons why our SDEB classification do not achieve the expected quality parameters, can be attributed to the very limited sample size – only 154 SDEBs in total, of which just 106 were used for training and testing – posing a challenge for the model to learn reliable patterns. Unlike CEBs or DEBs, SDEB chi-square plots often lacked clear distinguishing features: many exhibited distorted damping behaviour, as shown in Figs 2(d) and (e), while others resembled CEB-like curves or the examples shown




Figure 6. The confusion matrix of the 1D-CNN model using chi-square plots from *Kepler* data is shown here. (a) Displays the confusion matrix for all three EB classes: CEB, DEB, and SDEB. (b) Displays the confusion matrix for the binary classification involving only the two major classes: CEB and DEB.

Table 5. Precision, Recall, and F1-scores for the classification of CEBs, DEBs, and SDEBs using chi-square plots from the *Kepler* and *Kepler + PHOEBE* data sets. CNN1 represents the experiment conducted with only *Kepler* data, while CNN2 represents the experiment conducted with the combined *Kepler* and *PHOEBE* data. The first three rows show the results for the C, D, and SD classes, and the following three rows show the results for only CEBs and DEBs.

Metric	C (CNN1)	D (CNN1)	SD (CNN1)	C (CNN2)	D (CNN2)	SD (CNN2)
Precision	0.84	0.97	0.49	0.98	0.98	0.96
Recall	0.91	0.94	0.47	0.99	0.99	0.77
F1-score	0.87	0.95	0.48	0.98	0.99	0.86
Precision	0.88	0.98	...	0.99	0.98	...
Recall	0.96	0.94	...	0.98	0.99	...
F1-score	0.92	0.96	...	0.98	0.99	...

in Fig. 2(f). Consequently, our SDEB classification rates remained low. To address these issues, we augmented our data set by combining actual *Kepler* light curves with synthetic light curves generated using PHOEBE modelling code (A. Prsa et al. 2011).

4.2.2 CNN using Kepler and PHOEBE data combined

PHOEBE (A. Prša et al. 2016) is an EB modelling code based on the Wilson–Devinney model (R. E. Wilson & E. J. Devinney 1971). Using the PHOEBE forward-modelling algorithm,⁴ we generated synthetic LCs along with 3D visualizations of each system’s surface geometry and orbital configuration (mesh plots) for DEB, SDEB, and CEBs. From these synthetic LCs, we then produced the corresponding chi-square versus box-size plots. To construct the models, we adopted appropriate PHOEBE configurations for D, SD, and CEB systems and varied key physical parameters such as mass ratio, P_{orb} , effective temperatures of both components, orbital inclination, and orbital

⁴We used PHOEBE version 2.4.15.




Figure 7. PHOEBE mesh, phase-folded light curve, and chi-square plots for D (a, b, c panels), SD (d, e, f panels), and C (g, h, i panels) binaries. All the data were generated from the CALEB catalogue.

separation. To ensure that the simulated chi-square plots were based on physically realistic systems, we used representative systems from the Catalogue and Atlas of Eclipsing Binaries (CALEB; D. H. Bradstreet 2004) as templates. The specific CALEB systems used were as follows: for DEBs, V541 Cyg (15.337 d), LV Her (18.435 d), and AI Phe (24.592 d); for SDEBs, DD Mon (0.568 d), XZ CMi (0.579 d), and BO Peg (0.58 d); and for CEBs, CC Com (0.2207 d), V523 Cas (0.2337 d), and BS Vul (0.4759 d). These systems were selected because their P_{orb} values were consistent with the range observed in our *Kepler* sample. From these representative systems, we generated a large number of synthetic LCs by systematically varying the binary parameters across physically plausible ranges. The resulting set of simulated LCs was then used to generate synthetic chi-square plots following the same procedure applied to the *Kepler* data. This simulated data set was subsequently used to train our 1D-CNN model.

In Fig. 7, examples of mesh plots, phase-folded LCs, and chi-square plots of D, SD, and CEBs are presented, which were generated using PHOEBE modelling code. In this figure, the DEB was based on the orbital parameters of V541 Cyg, the SDEB was based on XZ Pup, and the CEB was based on the V523 Cas system. The chi-square diagram in Fig. 7(c) differs from Figs 2(a), (b), and (c). However, it is expected as DEBs exhibited various chi-square shapes and in general they are not as distinct as the damping sinusoidal pattern observed in CEBs. Nevertheless, in all cases, the D chi-square plots displayed an exponentially increasing trend. The chi-square plots generated from PHOEBE data, compared to those generated with data from *Kepler*, revealed significant similarities for C and D systems. However, the SD chi-square plots from PHOEBE did not show the same variations seen in observed SDs and matched only a few cases.

We adopted a sample size of 13 872, consisting of both synthetic chi-square plots generated by PHOEBE and those from *Kepler* data (only with the highest epochs), which 2806 plots were from observed LCs and the remainder were from simulated data. Among the 6171 CEBs, 832 were from *Kepler* data. We adopted 7077 DEBs and 624 SDEBs, from which 1820 and 154 were from *Kepler* data, respectively.




Figure 8. The confusion matrix of the 1D-CNN model using chi-square plots from the combined *Kepler* and PHOEBE data is shown here. Panel (a) shows the confusion matrix for all three eclipsing binary classes: CEB, DEB, and SDEB. Panel (b) shows the confusion matrix for the binary classification involving only the two major classes: CEBs and DEBs.

We used a 60:40 split for training and testing the data. Here, the data was split using Scikit-learn's (F. Pedregosa et al.) built-in train test function, which performs random sampling based on the specified ratio. From the initial 60 per cent training data, we further set aside 20 per cent for validation, which resulted in 48 per cent training and 12 per cent validation samples, making the final data set consist of 48 per cent training and 12 per cent validation data and 40 per cent test data. The reason for using a higher proportion of data for testing was that our sample contained a large amount of synthetic data generated with PHOEBE, while the number of *Kepler* SD sample was very limited. To ensure that some *Kepler* data were included in the test set amid the abundance of synthetic data, we opted for a larger test set than usual. After testing our 1D-CNN model, we found that it had an increased overall accuracy of 98 per cent. Out of 2831 DEBs in the test set, 2804 were well classified, resulting in an accuracy of 99 per cent. For the 2468 CEBs in the test sample, 2433 were well classified, with an accuracy of 98.6 per cent. Finally, out of 250 SD targets in the test sample, 193 were well classified, resulting in an accuracy of 77.2 per cent, a clear improvement in the identification of this type. The resulting confusion matrix was shown in Fig. 8(a).

We also performed the same analysis with SDEBs excluded, keeping the train, validation, and test sample ratio the same as before, at 48:12:40. In this case, the model was able to distinguish between C and D classes with 99 per cent accuracy. Out of 2466 CEBs in the test sample, 2433 were well classified, resulting in an accuracy of 98.7 per cent. Additionally, out of 2826 D instances in the test sample, 2804 were well classified, resulting in an accuracy of 99.2 per cent. The confusion matrix for this analysis was shown in Fig. 8(b). In Table 5, we presented the Precision, Recall, and F1-scores for these two analyses. The columns labelled C (CNN2), D (CNN2), and SD (CNN2) represented the metrics computed using *Kepler* and PHOEBE data combined. The first and last three rows corresponded to CNN2 with three classes and with only CEBs and DEBs, respectively.

The classification results from both CNN methods (with and without simulated data) are also presented in Table 3. While the PDS method analysed all available targets, the two CNN methods had different test sets due to their independent and unbiased sampling strategies. Therefore, only the common targets among all three approaches are displayed to evaluate classification accuracy in Table 3⁵.

⁵The complete version of this table is available online.




Figure 9. Box plot of P_{PDS} values and the EB morphological classification. Each box plot is labelled on the horizontal axis, with correctly classified C, SD, and D systems shown as solid boxes and the misclassified C, SD, and DEBs shown with a hatch pattern. The horizontal line inside each box indicates their median value (Q2). Dots beyond the whiskers denote outliers.

5 DISCUSSION

5.1 Comparison between classifications

We compared our visual classification with the classification scheme in KEBC, that is based on the c value. If the EBs had a c value less than 0.1, they were classified as well D systems. c values between 0.1 and 0.5 indicated primarily DEBs, while SDEBs had c values between 0.5 and 0.7. CEBs had c values between 0.7 and 0.8, and ellipsoidal EBs had c values greater than 0.8 up to 1. We classified both CEBs and ellipsoidal EBs under the CEB category. Out of 2806 common EBs (excluding 59 UNC cases) between the KEBC and our sample, 2502 matched our visual classification, yielding an overall accuracy of 89 per cent with respect to the KEBC.

Additionally, comparing KEBC's classification with our method based on PDS, 2230 EBs matched both classification, representing 79.5 per cent of the sample.

A total of 2201 targets matched across automated, visual, and KEBC classification, resulting in an overall agreement of 78.5 per cent. The mismatch in classifications was primarily due to the SD and CEB overlapped classes. We observed that some targets, such as KIC 11230837 and KIC 8288741, had c values of 0.61 and 0.64, respectively, which were classified as SD according to the KEBC catalogue, but visually they appeared to be DEBs. Moreover, there was ambiguity in cases where the c value was exactly 0.7, as both SD and CEB classifications overlapped at this threshold.

The PDS function showed limitations in classifying high- P_{orb} CEBs and SDEBs (Fig. 9). Compared to this, the 1D-CNN (Section 4.2.1) improved classification performance, particularly for SDEBs. In the second analysis (CNN2), augmenting with PHOEBE-generated synthetic data yielded 99 per cent accuracy for CEB and DEB systems. However, CNN2 reflected bias introduced by the synthetic data when classifying the three classes, which was primarily observed in SDEBs, as the synthetic data did not capture the physical origin of the distorted nature in the chi-square plots. Thus, while CNN2 excelled at separating CEBs and DEBs, CNN1 provided the most reliable overall classification across our different methods.

Earlier studies employing machine learning models classified EBs primarily based on the morphology of their LCs. M. Čokina et al. (2021) trained their deep-learning model on 491 425 synthetic LCs generated using the ELISA software and evaluated it on 100 observed LCs from published sources. Their model was trained on overcontact and DEBs, while the evaluation set also contained SDEBs that were treated as DEBs during testing. The model achieved an overall accuracy of 98 per cent which increased to 100 per cent when the SDEBs were excluded from the evaluation set. In contrast, our work classified *Kepler* EB data into three primary classes (DEBs, SDEBs, and CEBs), achieving an SDEB accuracy of 47 per cent, and 99 per cent accuracy when classification was performed only on DEBs and CEBs using both synthetic and *Kepler* data combined.

B. Ulaş (2025) reported an overall accuracy of 92 per cent, with class-wise F1 scores of 0.935 (CEB), 0.940 (DEB), and 0.872 (SDEB). Their data set combined *Kepler*, ASAS, and CALEB light curves, using catalogue-based morphology labels and a balanced sample across the three classes. In comparison, our chi-square-based CNN achieved an overall accuracy of 90.3 per cent. The model produced similar performance for CEBs and DEBs systems, comparable to LC based approaches, while the SD class showed lower accuracy. The lower SDEB accuracy was influenced by intrinsic morphological complexity and class imbalance, as SD systems formed a much smaller fraction of our data set. In addition, we analysed the full *Kepler* catalogue which increased the classification difficulty. The aim of this work was not to replace LC-based classification, but to assess whether chi-square morphology could provide class-wise discrimination and uncover variability signatures beyond static LC shapes (Section 5.3).




Figure 10. Distribution of EB targets in the parameter space defined by P_{PDS} and P_{orb} . Panels show the three EB classes separately: (a) DEBs, (b) SDEBs, and (c) CEBs. Targets for which the visual classification agrees with the PDS classification are shown as solid blue markers, while targets where the two classifications disagree are represented by open black circle markers.

5.2 PDS period correlation with orbital period

Analysing the different P_{PDS} values in Table 3, we observed that those for some DEBs, or systems with DEB-like chi-squared plots, were generally higher than those for CEBs, and in some cases, unusually large values were observed. For example, KIC 2987433 had a P_{PDS} value of 1.37×10^7 box-size steps, which was extremely large, yet it was classified as a DEB according to the fitting criteria in Table 2. This occurred due to the non-sinusoidal shape of the chi-squared plot for DEBs, which affected the PDS fitting. To further support this observation, we examined the visually classified SDEB systems such as KIC 6046061 and KIC 6516874 and they also had extremely large P_{PDS} values and were misclassified by the PDS function as DEBs, again due to their non-sinusoidal chi-squared shapes. This confirmed that the extreme values arose from the PDS function's limitations in accurately fitting such profiles.

The results of our PDS classification scheme using box plots for the correctly assigned C (black), SD (orange), and D (blue) systems are presented in Fig. 9, showing a clear separation between them in terms of P_{PDS} . The misclassified EBs (hatch pattern) are also illustrated, displaying relatively higher and lower P_{PDS} values in respect to their classes. Each box represents the interquartile range (IQR), containing the middle 50 per cent of the data from the first quartile (Q1) to the third quartile (Q3). The IQR is defined as $\text{IQR} = Q3 - Q1$, and it provides statistical dispersion by minimizing the influence of extreme values. The green lines inside the boxes indicate the median P_{PDS} values (Q2), while whiskers extend from Q1 to the lowest value within $Q1 - 1.5 \times \text{IQR}$ and from Q3 to the highest value within $Q3 + 1.5 \times \text{IQR}$. These whiskers illustrate the typical range of variability in the P_{PDS} values for each classification, excluding extreme outliers. Data points beyond these whiskers represent outliers, highlighting extreme values within each category. We omitted a few large P_{PDS} values from Fig. 9 for better visualization, and the y-axis P_{PDS} values were plotted on a logarithmic scale. The systems that presented a good classification had well-separated PDS periods.

The misclassified EBs had P_{PDS} values outside the expected range for their class, except for SD systems. For CEBs, the median P_{PDS} was 8.88 box-size steps, whereas for misclassified C targets, this value was comparatively higher at 51.04 box-size steps, which was far from the threshold ($R^2 \geq 0.9$ and $P_{\text{PDS}} < 30$ or $P_{\text{PDS}} \leq 10$), established for the classification (see Table 2). For DEBs, the median P_{PDS} was 180.62 box-size steps. However, the misclassified DEBs showed a lower value of only 24.79 box-size steps, also outside our empirically defined threshold ($R^2 \geq 0.9$ and $P_{\text{PDS}} \geq 30$ or $P_{\text{PDS}} \geq 50$). The misclassified SDEBs did not present a clear P_{PDS} value separation like that seen for DEBs or CEBs. In fact, the SD P_{PDS} values showed an alignment with comparatively larger P_{PDS} values of C and smaller P_{PDS} values of DEBs, as seen in Fig. 9, where for the well-classified SD systems the median value was 22.15 box-size steps in comparison to 29.32 box-size steps for misclassified SDs. This could be in part because the classification scheme adopted was primarily based on the R^2 value.

In Fig. 10, we explored the correlation between P_{orb} and the corresponding P_{PDS} values for different EB morphological classes. The three panels show the distributions separately for (a) DEBs, (b) SDEBs, and (c) CEBs. In each panel, correctly classified systems are shown as solid blue markers, while misclassified systems are represented by open black circle markers. We considered systems with P_{PDS} values up to 100 box-size steps and P_{orb} less than 5 d. This range included most CEBs and SDEBs as well as short-period DEBs, which allowed us to better identify trends and examine the relationship between the orbital period and the corresponding P_{PDS} values across the different EB classes.

The majority of EB systems in our sample with P_{orb} of less than 5 d seemed to follow two different tracks: one where most of the CEBs lay and another where we found only D and SD systems. Most misclassified CEBs tended to have larger P_{orb} , while some systems with smaller P_{orb} did not follow the main C track, as shown in Fig. 10(c). These targets may have had ambiguous visual classifications, as they did not show clear LC behaviour, which could have led to a misclassification. Thus, the P_{PDS} and P_{orb} comparison plot helped to detect such cases. From Fig. 10(a), we observed that misclassified DEBs generally had shorter P_{orb} values, which could be a result of the previously defined criteria shown in Table 2. However, misclassified SDEBs may have chi-square plots either like C or D systems,

and the majority of them is often aligned with the C track, as shown in Fig. 10 (b). By examining the period comparison plot, we could identify these ambiguous systems, which contributed to the drop in accuracy of our classification method.

We developed an automated technique to identify the trends shown in Fig. 10. The algorithm first calculated the ratio $P_{\text{PDS}}/P_{\text{orb}}$ for each data point. It then applied the unsupervised machine learning model, k-means clustering (J. MacQueen 1967), on these ratios to divide the data into two distinct trend groups. For both clusters, the code performed independent linear regressions to determine the best-fitting equations. 10 percent of the outliers were removed within each cluster to ensure that the slopes were not influenced by them.

We identified that 1051 EBs were used by our automated technique to determine the first slope (solid red line in Fig. 13), and the equation of this slope is

$$P_{\text{orb}} = 0.0415P_{\text{PDS}} - 0.001. \quad (2)$$

The Pearson correlation coefficient for the first slope was 0.99, where the majority of the sample, 727 out of 1051, were CEBs. The code used 148 EBs to determine the second slope (dashed red line in Fig. 13), where the majority of the sample consisted of DEBs and SDEBs (121 DEBs and 21 SDEBs). This second slope is of the form

$$P_{\text{orb}} = 0.02P_{\text{PDS}} + 0.006, \quad (3)$$

with a Pearson correlation coefficient of 0.76.

We note that the slopes in equations (2) and (3) are specific to the *Kepler* long-cadence sampling interval of 29.43 min (1765.5 s). Since P_{PDS} is measured in box-size steps, each corresponding to one long-cadence epoch, these relations would yield different slopes if applied to data with a different cadence, such as *TESS* or *Kepler* short-cadence observations.

5.3 Temporally varying systems

Considering the *Kepler* data, the majority of the sample showed consistent chi-square plots across different quarters, and consequently similar phase-folded LCs. However, a subset displayed clear quarter-to-quarter changes in the chi-square plots across, followed by corresponding changes in their LCs. We refer to these objects as temporally varying (TV) systems.

Fig. 11 illustrated two representative SDEBs: a TV system (KIC 2577756) and a stable system (KIC 10920314). For KIC 2577756, both the chi-square plots and phase-folded LC changed between quarters, indicating a temporal evolution in the LC shape. In contrast, KIC 10920314 showed consistent chi-square structure and an essentially unchanged phase-folded LC across quarters. Similar quarter-scale LC variations in *Kepler* EBs were well known and were commonly associated with magnetic activity (e.g. evolving starspots), flares, and O'Connell-effect variability (M. F. Knote et al. 2022).

To identify TV candidates, we measured the P_{PDS} independently for all available quarters for each EB. Since TV systems exhibited changing chi-square morphologies across quarters, their fitted P_{PDS} values tended to vary more strongly than in stable systems. We therefore defined ΔP_{PDS} as the difference between the maximum and minimum P_{PDS} values measured across all available quarters, and adopted the normalized spread $\Delta P_{\text{PDS}}/P_{\text{PDS}}$ as a proxy for temporal variability. To avoid cases where the PDS fit became unreliable, we discarded targets with $P_{\text{PDS}} > 100$ box-size steps, and we restricted the analysis to systems where the PDS fit was well defined (primarily CEBs, SDEBs, and short-period DEBs). We also excluded long-period DEBs, where polynomial-like chi-square trends could dominate and affect the reliability of the derived P_{PDS} values; specifically, we adopted $P_{\text{orb}} \leq 4.2$ d, consistent with $P_{\text{PDS}} = 100$ in equation (2).

To calibrate a practical threshold for TV selection, we cross-matched our sample with the catalogue of TV systems compiled by M. F. Knote et al. (2022), identifying 72 systems in common (hereafter the test subset). We computed $\Delta P_{\text{PDS}}/P_{\text{PDS}}$ for this subset and examined its distribution (Fig. 12). The resulting distribution was strongly right-skewed, as confirmed by a Weibull-based skewness analysis (W. Weibull 1951), and we therefore adopted a Median–MAD criterion to define a robust threshold. This yielded a threshold value of 0.0133, and systems with $\Delta P_{\text{PDS}}/P_{\text{PDS}}$ above this value were classified as TV candidates.

Applying this threshold to our primary sample selected 416 TV systems. Fig. 13 places these objects in the $P_{\text{PDS}}-P_{\text{orb}}$ plane. For each EB, the minimum and maximum quarter-wise P_{PDS} values are connected by a grey segment to visualize the temporal spread. Most CEBs cluster tightly along the main correlation (solid red line), whereas many DEBs, and some SDEBs, exhibit larger spreads and tend to lie closer to the secondary trend (dashed red line). A small number of systems deviate from the main locus; in such cases, the chi-square morphology departs from the sinusoidal pattern, which can shift the fitted P_{PDS} values while still producing large $\Delta P_{\text{PDS}}/P_{\text{PDS}}$. Representative examples include KIC 5302006, a short-period system with a nearly sinusoidal light curve that was visually classified as a CEB but occupies the lower right region of Fig. 13. According to G. Matijević et al. (2012), such systems have morphology parameters in the range $0.8 \leq c \leq 1.0$, and in the KEBC (B. Kirk et al. 2016) KIC 5302006 has $c = 0.97$, consistent with an ellipsoidal EB. Another example is KIC 8432040, which has $P_{\text{orb}} = 5.3$ d and $c = 0.97$ in the KEBC, but exhibits polynomial-like chi-square trends and was classified as a DEB by the PDS function. We found that approximately 85 per cent of unmatched CEBs from the PDS classification are ellipsoidal systems, with most having $c \geq 0.9$. While these cases represent edge effects of the chi-square morphology, they do not undermine the overall effectiveness of the $\Delta P_{\text{PDS}}/P_{\text{PDS}}$ metric in identifying TV systems.

Finally, Table 6 lists a representative subset of our TV classifications, including four EBs that showed no literature references to any interesting physical characteristics such as star spots, flares, or the O'Connell effect, which could lead to TV behaviour. These systems might represent newly identified candidates that require further investigation in future studies.




Figure 11. A varying SD system, KIC 2577756, and a stable SD system, KIC 10920314, are shown here. Panel (a) illustrates the chi-square plot of quarter 9 for KIC 2577756, and Panel (b) shows its phase-folded LC for the same quarter. Panel (c) presents the chi-square plot of quarter 14 for the same system, while Panel (d) shows the corresponding phase-folded LC. Panel (e) displays the chi-square plot of quarter 9 for the stable SD system KIC 10920314, and Panel (f) shows its LC for the same quarter. Panel (g) illustrates the chi-square plot of quarter 15 for KIC 10920314, and panel (h) shows the corresponding LC.




Figure 12. Distribution of $\Delta P_{\text{PDS}}/P_{\text{PDS}}$ for the test subset. The x-axis represents $\Delta P_{\text{PDS}}/P_{\text{PDS}}$, shown with logarithmically spaced bins. The red dashed line marks the median of the distribution, while the black dashed line indicates the median–MAD threshold. The solid blue and green vertical lines correspond to KIC 2577756 ($\Delta P_{\text{PDS}}/P_{\text{PDS}} = 1.234$) and KIC 7284688 ($\Delta P_{\text{PDS}}/P_{\text{PDS}} = 2.612$), respectively.




Figure 13. Distribution of TV systems with P_{PDS} on the x-axis and P_{orb} on the y-axis. The solid red line shows equation (2), while the dashed red line shows equation (3). For each system, quarter-wise P_{PDS} values were connected with grey lines to indicate the spread across quarters. Visually classified CEBs are shown in black, DEBs in blue, and SDEBs in orange. The minima were shown as hollow circles, while the maxima were shown as solid circles in their corresponding EB colors.

5.4 The TV nature

To investigate the origin of the TV behaviour, we cross-matched our TV sample with 416 systems reported in relevant *Kepler* EB studies using *Kepler* IDs and system coordinates. Using the catalogue of stellar temperatures for *Kepler* EB (D. J. Armstrong et al. 2013), we found that cooler stars (late-F, G, K, and M types) exhibited systematically larger $\Delta P_{\text{PDS}}/P_{\text{PDS}}$ values than hotter stars. This trend is consistent with enhanced magnetic activity in stars with convective envelopes, driven by rotational dynamos, leading to increased starspot coverage (K. G. Strassmeier 2009), higher flare rates – particularly in K and M dwarfs (J. R. A. Davenport 2016) – and coronal X-ray emission from magnetically confined plasma (G. S. Vaiana et al. 1981; N. J. Wright et al. 2011).

To further constrain the physical origin of the observed variability, we cross-matched the 416 TV systems with several catalogues related to starspot modulation, flares, and O’Connell-effect variability (J. Deboscher et al. 2011; B. Balaji et al. 2015; J. R. A. Davenport 2016; Q. Gao et al. 2016; M. F. Knotte et al. 2022; J. C. Lurie et al. 2017). The cross-matched systems and their reported properties are summarized in Table 6, where SM denotes starspot modulation, F and S denote flares and starspots, OE denotes the O’Connell effect, and AM denotes asymmetric minima.

From this cross-match, we found a median $\Delta P_{\text{PDS}}/P_{\text{PDS}}$ value of 0.5 for flare-active close binaries (Q. Gao et al. 2016), which is significantly higher than the adopted threshold of 0.0133. As an illustrative example, Fig. 14(a) shows KIC 2577756, which has

Table 6. In this table, we present a small subset of our TV classifications. For each target, we listed the *Kepler* ID, visual classification, P_{PDS} , P_{orb} , and $\Delta P_{\text{PDS}}/P_{\text{PDS}}$ values, along with references from the literature survey and the physical characteristics of each EB in the Notes column. The last four systems, for which no references are provided, are new TV systems identified using our method. The full version of this table is available online.

<i>Kepler</i> ID	Visual classification	P_{PDS}	P_{orb}	$\Delta P_{\text{PDS}}/P_{\text{PDS}}$	Reference ^a	Notes ^b
KIC 2305372	S,D	32.274	1.405	1.038	[2,3,5,6]	F,OE
KIC 2305543	D	64.273	1.362	0.037	[2]	SM
KIC 2309587	D	87.269	1.839	0.055	[2,3,5]	SM,F
KIC 2447893	S,D	16.068	0.662	0.994	[1,2,3,4]	SM,F
KIC 2557430	D	78.001	1.298	0.471	[1,2,3,5]	SM,F
KIC 2558370	D	58.891	0.546	0.172	[2]	SM
KIC 2569494	S,D	67.025	1.523	0.097	[2,6]	SM,AM,OE
KIC 2577756	S,D	20.478	0.870	1.234	[2,3,5,6]	SM,F,AM,OE
KIC 2695740	C	85.968	3.616	0.096	[1,2,3,4,6]	SM,F,S,OE
KIC 2720354	D	66.030	2.821	0.049	[1,2,3]	SM,F
KIC 5428668	D	30.532	1.406	0.095	...	...
KIC 8800998	C	20.9	0.883	0.013	...	...
KIC 9077796	C	11.566	0.238	0.503	...	...
KIC 10991769	C	76.405	3.281	0.02	...	...

^a Reference codes: 1: J. R. A. Davenport (2016), 2: J. C. Lurie et al. (2017), 3: J. Debosscher et al. (2011), 4: B. Balaji et al. (2015), 5: Q. Gao et al. (2016), 6: M. F. Knote et al. (2022).

^b SM : literature indicates the presence of Starspot Modulation; F : literature indicates the presence of Flare; AM : literature indicates the presence of Asymmetric Minima; OE : literature indicates the presence of O’Connell Effect; S : literature indicates the presence of Spots




Figure 14. (a) Time-domain *Kepler* light curve of KIC 2577756, plotted as normalized flux versus BJD, showing strong variability and clear flare signatures consistent with its high $\Delta P_{\text{PDS}}/P_{\text{PDS}}$ value. (b) Time-domain *Kepler* light curve of KIC 7284688, plotted as normalized flux versus BJD, showing strong variability and varying O’Connell effect due to asynchronous rotation of starspots consistent with its high $\Delta P_{\text{PDS}}/P_{\text{PDS}}$ value.

$\Delta P_{\text{PDS}}/P_{\text{PDS}} = 1.234$ and exhibits strong flare activity in its time-domain light curve, consistent with previous studies. This system is also highlighted in Fig. 12 by a blue vertical line.

Fig. 14(b) presents KIC 7284688, a solar-type EB displaying a rapidly varying O’Connell effect attributed to asynchronous starspot rotation (Y. Pan & X. Zhang 2023). This system shows $\Delta P_{\text{PDS}}/P_{\text{PDS}} = 2.612$ and is marked by a green vertical line in Fig. 12.

Overall, these results reinforce the strong connection between magnetic activity and temporal variability in our TV sample, and demonstrate that flare- and starspot-active systems systematically exhibit higher $\Delta P_{\text{PDS}}/P_{\text{PDS}}$ values, making them particularly important targets for detailed follow-up studies.

6 SUMMARY AND CONCLUSIONS

In this study, we classified 2806 EBs into CEBs, DEBs, and SDEBs using long-cadence *Kepler* data. We created a new classification scheme that was different from the traditional LC shape analysis, which was based on an indirect measure of the time-series data and that did not require any LC modelling. We generated LCs from MAST FITS files cross-matched with the KEBC and normalized them via quarter-wise chi-square versus box-size analysis, which revealed characteristic trends for each EB type. We used box sizes between 10 and 101, which showed the most distinctive chi-square versus box-size shapes when tested on a few targets and were also chosen for practical computational reasons. Because the chi-square patterns were stable across different quarters for most EBs, we generated plots

for all available quarters to provide sufficient training data for the network, while reserving the highest epoch quarter of each system exclusively for the testing set.

In this analysis, we first fitted a PDS function to the quarter with the largest number of epochs, extracted the P_{PDS} values, and, using the conditions on R^2 and P_{PDS} , achieved a classification accuracy of 86.5 per cent. In addition, for CEBs, we found a strong linear correlation between P_{PDS} and P_{orb} , with a Pearson coefficient of 0.99, confirming that P_{PDS} is intrinsically linked to the P_{orb} of the EBs.

To improve the PDS methodology, we built a 1D-CNN for the morphological classification and conducted two different analysis to identify the best-performing model and data set. We found that chi-square plots remained consistent across quarters for most CEBs and DEBs, while some SDEBs and low- P_{orb} DEBs showed inconsistencies. By using data from all available quarters and removing targets with inconsistent chi-square patterns, we improved the model's ability to learn general features and achieved an accuracy of 90 per cent. In a separate experiment using only CEBs and DEBs, our method distinguished these two classes with an overall accuracy of 95 per cent.

To further enhance performance, we combined observed chi-square plots with synthetic data generated using PHOEBE. This approach improved classification, particularly for CEBs and DEBs. However, PHOEBE models still struggled to reproduce the full variability observed in SDEBs from *Kepler* LC data. This resulted in more homogeneous chi-square plots and contributed to the improved accuracy for SDEBs. When we tested the PHOEBE model only on CEBs and DEBs, the accuracy increased to 99 per cent, since chi-square variations in these classes were primarily due to changes in P_{orb} , which PHOEBE replicated effectively by adjusting the period.

The CNN1 provided the best accuracy across the three EB classes together, correctly representing the observed data. In contrast, the CNN2 achieved the highest accuracy when classifying only CEBs and DEBs but was unable to classify the SDEBs properly, as we believed the variation in chi-square plots for SDEBs was linked to various physical phenomena.

To investigate the misclassified systems, we identified EBs whose chi-square plots and light curves varied from quarter to quarter. We termed these systems as TV systems and developed an automated method to detect them. TV systems exhibited quarterly chi-square variations that led to changes in P_{PDS} values. By measuring the difference between the maximum and minimum P_{PDS} values across quarters, we quantified the extent of variability, which was also reflected in their LCs.

The chi-square plots were very sensitive to variations in the binary LC, especially those generated due to magnetic activity, which was one of the prime reasons for temporal variation in EBs. Our method detected even small changes in chi-square plots, as their shapes varied from quarter to quarter. After normalizing these variations, we obtained the $\Delta P_{\text{PDS}}/P_{\text{PDS}}$ parameter, which detected systems with enhanced magnetic activity. This sensitivity enabled us to identify four EBs without any literature references, which require further investigation.

The chi-square plots not only helped in classifying the EBs but were also highly sensitive to physical changes in the binaries. We observed that DEBs with shorter P_{orb} tended to show damping-like chi-square patterns, though not as strongly as CEBs. In contrast, CEBs exhibited more stable, damping-like plots, possibly because these systems had reached equilibrium through mass transfer or other processes. SDEBs did not show very distinctive chi-square patterns, as most of the SD sample was dominated by Temporally Varying systems, showing variable and distorted chi-square shapes. Leveraging these variations, we classified the targets using the PDS function and CNN. In future work, we plan to explore these unique chi-square shapes further and relate them to specific physical phenomena, which may lead to a more accurate classification of SDEBs and help unravel the magnetic effects influencing the chi-square plots.

ACKNOWLEDGEMENTS

We are very grateful to the anonymous referee for their illuminating and insightful suggestions, which led to substantial improvements in the manuscript.

PC, JA acknowledges support from the Spanish Virtual Observatory (<https://svo.cab.inta-csic.es>) project funded by the Spanish Ministry of Science, Innovation and Universities /State Agency of Research MCIU/AEI/10.13039/501100011033 through grant PID2023-146210NB-I00. HJ acknowledges Science and Technology Facilities Council (STFC) support for this research provided through infrastructure grant ST/R000905/1. MCGO acknowledges financial support from the Agencia Estatal de Investigación (AEI/10.13039/501100011033) of the Ministerio de Ciencia e Innovación and the European Regional Development Fund (ERDF) 'A way of making Europe' through project PID2022-137241NB-C42. JA's contribution to this work is part of his academic activities as a professor at Universidad Militar Nueva Granada, Bogotá, Colombia.

CONFLICT OF INTEREST

The authors declare that there are no conflicts of interest regarding the publication of this paper.

DATA AVAILABILITY

All code developed for this work is available at the following GitHub repository: <https://github.com/mousam100>

The full machine-readable classification table and the complete temporally varying systems (TVS) table have been submitted to the CDS. These data are available at the CDS via anonymous ftp to cdsarc.u-strasbg.fr (130.79.128.5), or via <https://cdsarc.cds.unistra.fr/viz-bin/cat/J/other/RASTI>. Following the VizieR procedure, the catalogue will be released publicly after publication of this paper.

REFERENCES

- Andersen J., 1991, *A&ARv*, 3, 91
- Armstrong D. J., Gómez Maqueo Chew Y., Faedi F., Pollacco D., 2013, *MNRAS*, 437, 3473
- Balaji B., Croll B., Levine A. M., Rappaport S., 2015, *MNRAS*, 448, 429
- Bradstreet D. H., 2004, *Catalog and Atlas of Eclipsing Binaries*, Eastern University, St. Davids
- Breiman L., 2001, *Mach. Learn.*, 45, 5
- Čokina M., Maslej-Krešňáková V., Butka P., Parimucha Š., 2021, *Astron. Comput.*, 36, 100488
- Cruz P., Aguilar J. F., Garrido H. E., Diaz M. P., Solano E., 2022, *MNRAS*, 515, 1416
- Davenport J. R. A., 2016, *ApJ*, 829, 23
- Daza-Perilla I. V., Gramajo L. V., Lares M., Palma T., Ferreira Lopes C. E., Minniti D., Clariá J. J., 2023, *MNRAS*, 520, 828
- Debosscher J., Blomme J., Aerts C., De Ridder J., 2011, *A&A*, 529, A89
- Ding X., Ji K., Cheng Q., Song Z., Wang J., Tian X., Wang C., 2025, *AJ*, 169, 202
- Frost W. et al., 2024, *ApJ*, 972, 199
- Gao Q., Xin Y., Liu J.-F., Zhang X.-B., Gao S., 2016, *ApJS*, 224, 37
- Hempel M. et al., 2014, *Messenger*, 155, 24
- Kallrath J., Milone E. F., 2009, *Eclipsing Binary Stars: Modeling and Analysis*. Springer, New York
- Kingma D. P., Ba J., 2015, *International Conference on Learning Representations (ICLR)*. San Diego, CA
- Kirk B. et al., 2016, *AJ*, 151, 68
- Knote M. F., Caballero-Nieves S. M., Gokhale V., Johnston K. B., Perlman E. S., 2022, *ApJS*, 262, 10
- Li P., Bahri O., Filali Boubrahimi S., Hamdi S. M., 2025, *ApJS*, 280, 52
- Lin T.-Y., Goyal P., Girshick R., He K., Dollár P., 2020, *IEEE Trans. Pattern Anal. Mach. Intell.*, 42, 318
- Lurie J. C. et al., 2017, *AJ*, 154, 250
- MacQueen J., 1967, in Le Cam L. M., Neyman J. eds, *Proc. 5th Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 1. University of California Press, Berkeley, CA, p. 281
- Matijević G., Prša A., Orosz J. A., Welsh W. F., Bloemen S., Barclay T., 2012, *AJ*, 143, 123
- Minniti D. et al., 2010, *New Astron.*, 15, 433
- Minniti D. et al., 2018, *A&A*, 616, A26
- Morris R. L., Twicken J. D., Smith J. C., Clarke B. D., Jenkins J. M., Bryson S. T., Girouard F., Klaus T. C., 2020, in Jenkins J. M., ed., *Technical Report KSCI-19081-003*. NASA Ames Research Center, Moffett Field, CA
- Murphy S. J., 2012, *MNRAS*, 422, 665
- Pan Y., Zhang X., 2023, *AJ*, 165, 247
- Pedregosa F. et al., 2011, *J. Mach. Learn. Res.*, 12, 2825
- Prsa A., Matijevic G., Latkovic O., Vildardell F., Wils P., 2011, *Astrophysics Source Code Library*, record ascl:1106
- Prša A. et al., 2016, *ApJS*, 227, 29
- de Ridder S., Duin R., 2002, *Technical Report PH-2002-01*, Pattern Recognition Group. Delft University of Technology, Delft
- Roweis S. T., Saul L. K., 2000, *Science*, 290, 2323
- Saito R. K. et al., 2012, *A&A*, 537, A107
- Sarro L. M., Sánchez-Fernández C., Giménez Á., 2006, *A&A*, 446, 395
- Slawson R. W. et al., 2011, *AJ*, 142, 160
- Strassmeier K. G., 2009, *Astron. Astrophys. Rev.*, 17, 251
- Torres G., Andersen J., Giménez A., 2010, *A&ARv*, 18, 67
- Ulaş B., 2023, *Turkish Journal of Astronomy and Astrophysics*, 6, 1
- Vaiana G. S. et al., 1981, *ApJ*, 245, 163
- Van Cleve J. E. et al., 2016, *Kepler Data Characteristics Handbook*, Kepler Science Document KSCI-19040-005. NASA Ames Research Center, Moffett Field, CA
- Virtanen P. et al., 2020, *Nature Methods*, 17, 261
- Weibull W., 1951, *J. Appl. Mech.*, 18, 293
- Wilson R. E., Devinney E. J., 1971, *ApJ*, 166, 605
- Wright N. J., Drake J. J., Mamajek E. E., Henry G. W., 2011, *ApJ*, 743, 48

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.