



The TESS Ten Thousand Catalog: 10,001 Uniformly Vetted and Validated Eclipsing Binary Stars Detected in Full-frame Image Data by Machine Learning and Analyzed by Citizen Scientists

Veselin B. Kostov^{1,2}, Brian P. Powell¹, Aline U. Fornear³, Marco Z. Di Fraia³, Robert Gagliano⁴, Thomas L. Jacobs⁵, Julien S. de Lambilly³, Hugo A. Durantini Luca³, Steven R. Majewski⁶, Mark Omohundro⁷, Jerome Orosz⁸, Saul A. Rappaport⁹, Ryan Salik^{3,10}, Donald Short⁸, William Welsh⁸, Svetoslav Alexandrov^{3,11}, Cledison Marcos da Silva³, Erika Dunning⁸, Gerd Gühne³, Marc Hutten³, Michiharu Hyogo³, Davide Iannone³, Sam Lee³, Christian Magliano¹², Manya Sharma^{3,13}, Allan Tarr³, John Yablonsky³, Sovan Acharya³, Fred Adams¹⁴, Thomas Barclay¹, Benjamin T. Montet¹³, Susan Mullally¹⁵, Greg Olmschenk¹, Andrej Prša¹⁶, Elisa Quintana¹, Robert Wilson¹, Hasret Balcioğlu³, Ethan Kruse¹, and

The Eclipsing Binary Patrol Collaboration

¹ NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA

² SETI Institute, 189 Bernardo Ave., Ste. 200, Mountain View, CA 94043, USA

³ Citizen Scientist, Eclipsing Binary Patrol Collaboration

⁴ Citizen Scientist, Glendale, AZ 85308, USA

⁵ Citizen Scientist, Missouri City, TX 77459, USA

⁶ Department of Astronomy, University of Virginia, 530 McCormick Rd., Charlottesville, VA 22904, USA

⁷ Citizen Scientist, c/o Zooniverse, Department of Physics, University of Oxford, Denys Wilkinson Bldg., Keble Rd., Oxford OX1 3RH, UK

⁸ Department of Astronomy, San Diego State University, 5500 Campanile Dr., San Diego, CA 92182, USA

⁹ Department of Physics, Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

¹⁰ Department of Computer Science, Princeton University, NJ, USA

¹¹ Institute of Plant Physiology and Genetics, Bulgarian Academy of Sciences, Acad. G. Bontchev Str., Bl. 23, Sofia 113, Bulgaria

¹² Dipartimento di Fisica “Ettore Pancini,” Università di Napoli Federico II, 80126 Napoli, Italy

¹³ School of Physics, University of New South Wales, Sydney, NSW 2052, Australia

¹⁴ Department of Physics, University of Michigan, Ann Arbor, MI, USA

¹⁵ Space Telescope Science Institute, 3700 San Martin Dr., Baltimore MD 21212, USA

¹⁶ Department of Astrophysics and Planetary Science, Villanova University, 800 Lancaster Ave., Villanova, PA 19085, USA

Received 2025 May 8; revised 2025 June 5; accepted 2025 June 5; published 2025 August 1

Abstract

The Transiting Exoplanet Survey Satellite (TESS) has surveyed nearly the entire sky in full-frame image mode with a time resolution of 200 s to 30 minutes and a temporal baseline of at least 27 days. In addition to the primary goal of discovering new exoplanets, TESS is exceptionally capable at detecting variable stars, and in particular short-period eclipsing binaries, which are relatively common, making up a few percent of all stars, and represent powerful astrophysical laboratories for deep investigations of stellar formation and evolution. We combed Sectors 1–82 of the TESS full-frame image data searching for eclipsing binary stars using a neural network that identified ~ 1.2 million stars with eclipse-like features. Of these, we have performed an in-depth analysis on $\sim 60,000$ targets using automated methods and manual inspection by citizen scientists. Here we present a catalog of 10,001 uniformly vetted and validated eclipsing binary stars that passed all our ephemeris and photocenter tests, as well as complementary visual inspection. Of these, 7936 are new eclipsing binaries while the remaining 2065 are known systems for which we update the published ephemerides. We outline the detection and analysis of the targets, discuss the properties of the sample, and highlight potentially interesting systems. Finally, we also provide a list of $\sim 900,000$ unvetted and unvalidated targets for which the neural network found eclipse-like features with a score higher than 0.9, and for which there are no known eclipsing binaries within a sky-projected separation of a TESS pixel ($\approx 21''$).

Unified Astronomy Thesaurus concepts: Binary stars (154); Eclipsing binary stars (444)

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

1. Introduction

Binary stars make up a large fraction of the Galactic stellar population (e.g., D. Raghavan et al. 2010; A. Tokovinin 2021; S. S. R. Offner et al. 2023). Of these, perhaps the most important subsets are those that produce eclipses due to a favorable geometric configuration with respect to the observer.

These eclipsing binary stars (EBs) have paved the “royal road” to stellar astrophysics (H. N. Russell 1948) and have long served as a fundamental pillar on which our understanding of how stars form and evolve stands (e.g., D. E. Osterbrock 1953; J. Andersen 1991; G. Torres et al. 2010). Spectroscopic double-lined EBs enable direct and accurate measurements of the masses, radii, and temperatures of their components, and provide critical calibrators for theoretical models (e.g., G. Torres et al. 2010).

Despite the ubiquitous distribution of binary stars throughout the solar neighborhood and over two centuries of study



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

(e.g., J. Goodricke 1783; Z. Kopal 1956; O. J. Eggen 1957; V. Niemela 2001; and references therein, including Sewell’s letter to S. Vince) pressing questions about these systems remain. For example, it is unclear whether the multiplicity properties of stellar systems are universal or depend on the formation environment and/or stellar mass, what the origin is of the brown dwarf “desert” scarcity, and how stellar multiplicity affects planet formation (e.g., M. Moe & R. Di Stefano 2017 and references therein). These uncertainties are due in large part to the enormous size of the parameter space, since binary stars have extensive distributions of stellar masses and mass ratios, orbital periods, eccentricities, etc., all of which can vary with the environment (e.g., cluster membership).

Large-scale photometric surveys are well suited for monitoring a large number of binary stars through the detection of eclipses, and have detected hundreds of thousands of EBs. For example, millions of EBs have been observed by Gaia (N. Mowlavi et al. 2023);¹⁷ hundreds of thousands by the Optical Gravitational Lensing Experiment (OGLE; I. Soszyński et al. 2017), ASAS-SN (D. M. Rowan et al. 2022), ATLAS (A. N. Heinze et al. 2018), and the Wide-field Infrared Survey Explorer (WISE; E. Petrosky et al. 2021); and tens of thousands by primarily exoplanet-focused surveys such as Kepler (A. Prša et al. 2011; R. W. Slawson et al. 2011; K. E. Conroy et al. 2014a), the Transiting Exoplanet Survey Satellite (TESS; P. W. Sullivan et al. 2015), SuperWASP (H. B. Thiemann et al. 2021), etc. With its extremely wide sky coverage ($\sim 98\%$) and long dwell time (~ 27 days of nearly continuous observations), NASA’s TESS mission is an excellent example of the power of all-sky surveys for studying EBs. While the primary science objective of TESS is finding transiting rocky exoplanets around nearby stars (G. R. Ricker et al. 2015), it presents an ideal platform for the detection of thousands of EBs covering a wide range of physical and orbital parameter space (e.g., N. L. Eisner et al. 2021; L. Caciapuoti et al. 2022; E. L. Howard et al. 2022; A. Prša et al. 2022; M. J. Green et al. 2023; C. Magliano et al. 2023; M. Montalto 2023; L. W. IIspeert et al. 2024; E. J. Melton et al. 2024; X. Gao et al. 2025; Y. Shan et al. 2025).

The TESS mission is also well suited for exploring the variability of many different classes of stars, and searching for rare systems that may be studied with extensive follow-up observations from space and the ground. The large EB population monitored by TESS enables statistical studies of the effects of mass, mass ratio, and composition on the binary fraction, eccentricity, and orbital period distributions. Such a large sample of EBs covering all stellar types and Galactic environments also helps advance our knowledge of the physics of binary interactions, such as tidal forces, migration, spin-orbit coupling, and mass transfer (e.g., H. V. Zeipel 1910; Y. Kozai 1962; M. L. Lidov 1962; O. Pejcha et al. 2013; X. Fang et al. 2018; G. Fragione & B. Kocsis 2019; B. Liu & D. Lai 2019; A. S. Hamers et al. 2021; C. S. Kochanek 2021; M. M. Shara et al. 2021; A. A. Trani et al. 2022; P. Vynatheya & A. S. Hamers 2022; and references therein). Last but not least, by better understanding the distribution and properties of EBs in the Galaxy, we can improve our priors on background contamination for TESS’s core mission of exoplanet transit observations.

Given the enormous amount of data produced by the TESS mission—for example, there are, on average, ~ 3 million stars brighter than $T_{\text{mag}} = 15$ observed per sector—we need to develop sophisticated yet efficient analysis techniques to extract the relevant astrophysical information from this unique data set. At the time of writing, several projects have already developed pipelines for the extraction of full-frame image (FFI) lightcurves from TESS (e.g., R. J. Oelkers & K. G. Stassun 2018; A. D. Feinstein et al. 2019; D. A. Caldwell et al. 2020; M. Kunitomo et al. 2022; T. Han & T. D. Brandt 2023; J. D. Hartman et al. 2025), and have released tools and data products to the public. To study binary stars from TESS, we have developed a local implementation of the ELEANOR pipeline (A. D. Feinstein et al. 2019) and used it to extract FFI lightcurves for Sectors 1–82 for all targets brighter than $T_{\text{mag}} = 15$. To detect EB candidates, we have created and trained a machine learning (ML) identification scheme. Here we describe the development and implementation of our extraction and detection pipeline and the processing and analysis of the data by automated methods and human inspection and present the TESS Ten Thousand Catalog containing 10,001 uniformly vetted and validated EBs. Of these, 7936 are new EBs and 2065 are known EBs for which we update the ephemeris provided in one or more catalogs. We describe the general properties of the population and touch on individual systems of interest. The catalog provides general target information (TESS Input Catalog (TIC) ID, sky coordinates, TESS magnitude, number of sectors observed, effective temperature, and Gaia astrometric measurements), ephemerides, eclipse depths and durations, secondary phases, and relevant notes and comments. We envision this catalog as a community-facing product to serve as a platform for subsequent studies and analysis of both the population as a whole and individual targets of interest, including but not limited to confirmation and modeling efforts, crossmatching against catalogs of TESS planet candidates, etc. All our data products and results are publicly available as machine-readable online supplements.

This paper is organized as follows. In Section 2 we describe the construction of the FFI lightcurves; Section 3 outlines the identification of EB candidates by an ML pipeline while the vetting and validation of the candidates is presented in Section 4. Section 5 outlines the catalog of uniformly vetted and validated EBs, and the results are summarized in Section 6.

2. Construction of FFI Lightcurves

While other lightcurve data sets were available to us, such as the MIT Quick Look Pipeline (QLP; C. X. Huang et al. 2020a) or the TESS Science Processing Operations Center (SPOC) pipeline (D. A. Caldwell et al. 2020), we wanted to pursue potentially unknown systems beyond the scope of available lightcurves. For example, the QLP lightcurves are limited to stars brighter than $T_{\text{mag}} = 13.5$. As such, we undertook an effort to construct all available TESS lightcurves to a limit of $T_{\text{mag}} = 15.0$ using ELEANOR (A. D. Feinstein et al. 2019).¹⁸

We started by downloading the full TIC (K. G. Stassun et al. 2019), available as a set of CSV files in increments of 2°

¹⁷ Farewell, Gaia! Thank you for all the amazing science!

¹⁸ Later, we rebuilt these for public release using the ELEANOR-LITE pipeline for Sectors 1–26 to a limit of $T_{\text{mag}} = 16.0$ (B. P. Powell et al. 2022a). These are available at <https://archive.stsci.edu/hlsf/gsfsc-eleanor-lite>.

of decl., from MAST.¹⁹ Each target in the TIC was queried in parallel using the TESS-POINT Python package (C. J. Burke et al. 2020) to determine the sectors of TESS data in which it was present, effectively translating the overall TIC into a per-sector TIC.

In preparation for building the lightcurves for a given sector, we then downloaded the TESS FFIs from MAST and used ELEANOR to create the necessary “postcards” and “backgrounds” required for local construction of the lightcurves (described further in A. D. Feinstein et al. 2019). The per-sector TIC was then used as input to a parallelized implementation of ELEANOR on the NASA Center for Climate Simulation (NCCS) Discover supercomputer.²⁰ The outputs of our parallelized lightcurve construction code were minimized to limit the need for memory storage, and contained only basic metadata along with the times and fluxes.

3. ML Identification of EB Candidates

Eclipses are an ideal shape for ML classification in lightcurves. They are usually a prominent feature in the lightcurve, with common spatial interrelationships between the eclipse and the baseline, as well as the characteristic point at the eclipse minimum. These features are uniquely identifiable in lightcurves and lend themselves to processing with a convolutional neural network (CNN; Y. LeCun et al. 1989). Rather than limit ourselves to only those lightcurves that demonstrated periodicity with eclipses, we chose to pursue a strategy of training the neural network to find the feature of the eclipse. In this manner, we could also treat the lightcurve purely as a 1D shape rather than having to consider time dependencies, allowing a simpler methodology. Our intent was to build the neural network for classification purposes, i.e., to produce a single sigmoid-activated output where unity is a positive (indicating that the lightcurve contains an eclipse) and zero is a negative (indicating that the lightcurve does not contain an eclipse).

The performance of a CNN as a classifier is broadly tied to the depth of the network (K. Simonyan & A. Zisserman 2014). While vanishing or exploding gradients generally limit depth, the concept of residual blocks (K. He et al. 2016) has allowed for depth to be limited only by hardware (in terms of physical memory available) and by training data shape and batch size (in terms of the trade-off between depth and data size within the physical memory). As such, we designed the general structure of our neural network as a 1D adaptation of ResNet (K. He et al. 2016), which was originally designed to process 2D images. Of course, too much depth in conjunction with very little data or overly simplistic data can also prevent convergence. Since a lightcurve is not a particularly complex data representation requiring extreme depth, we started our development process with a relatively shallow network. We developed the neural network iteratively (using TensorFlow/Keras; M. Abadi et al. 2015; F. Chollet 2015), and made it deeper as we augmented our training data and ensured that additional data and a deeper network offered continued reduction of the model loss, as given by binary cross entropy (I. Goodfellow et al. 2016) using the RMSprop optimizer.²¹ We also found that an additive and multiplicative attention

mechanism (D. Bahdanau et al. 2014; M.-T. Luong et al. 2015) at the beginning of the network was beneficial to performance. Apart from the sigmoid activation on the output layer, we used leaky rectified linear unit (ReLU; V. Nair & G. E. Hinton 2010) activation throughout the network to prevent the problem of vanishing gradients. The structure of our neural network, shown in Figure 1, is rather simple. In total, our network has 241 layers with ~ 5.5 million trainable parameters.

3.1. Lightcurve Preprocessing

Concurrently with the development of the neural network, we also needed to refine our method of preprocessing the lightcurves, which, as in many ML applications, was critical to the performance of the neural network. ML methods require data to be of the same shape. This of course presents a problem with TESS FFI lightcurves, especially after the per-time-step quality flags are masked from the lightcurve, resulting in a wide variety of 1D array sizes. The temporal discontinuities caused by the data downlink gap and the quality mask also create a temporally irregular data set. We chose to ignore the temporal component and treat the lightcurve as a 1D shape rather than a time-dependent signal. This approach is consistent with our selection of a CNN rather than a recurrent neural network (D. E. Rumelhart et al. 1986) or another time-dependent methodology such as long short-term memory (LSTM; S. Hochreiter & J. Schmidhuber 1997), convolutional LSTM (ConvLSTM; X. Shi et al. 2015), or temporal convolutional networks (C. Lea et al. 2017), among others.

To create homogeneous shapes from the irregular lightcurves, our options were to either truncate longer lightcurves or pad shorter lightcurves. Truncation, of course, risks missing a lightcurve where a single eclipse occurs in the truncated section. Padding provides its own risks in providing artificial information in the discontinuity of the data shape. We decided to pad the lightcurves to a maximum length of 1400 elements, with the padding containing a mirror of the data. We emphasize again that the neural network has no time dependency, and therefore no understanding of periodicity, and it was determined that the neural network could learn to ignore the discontinuity in the padding in the same manner as it would the discontinuity in the collapsed time gaps. We also note that we developed the neural network during Year 2 of the TESS mission, when the 30 minute cadence data provided relatively short lightcurves. With the continued shortening of the TESS cadence over subsequent years, we have not retrained the network. Rather, we downsampled longer lightcurves to fit our required data shape.

TESS systematics presented a different set of challenges. In many TESS FFI lightcurves, indeed, the dominant signal is scattered light systematics, which often produce features in the lightcurves that resemble eclipses. In Years 1 and 2 of the TESS data, this problem is particularly pronounced in Sectors 1, 12, 13, 14, 15, 23, 24, and 26. In neural networks, and ML in general, the magnitude of a value is a representation of its importance. As such, strong eclipse-like systematic signals have the potential to dominate an ML method if not properly diminished in importance, hence the need for data scaling. We specifically selected the quantile transform as our method of scaling, using the built-in SCIKIT-LEARN (F. Pedregosa et al. 2011) package functionality. With this method, scattered light systematics are reduced to effectively the same size as the

¹⁹ https://archive.stsci.edu/tess/tic_ct1.html

²⁰ <https://www.nccs.nasa.gov/systems/discover>

²¹ <https://www.cs.toronto.edu/~hinton/coursera/lecture6/lec6e.mp4>




Figure 1. The structure of our neural network, with like layers grouped by color. The full network summary is shown on the left. The attention block has a structure such as shown on the upper right, while each convolutional block has a structure such as shown on the lower right. Arrows into the addition layers indicate the flow of the residual.

eclipses, forcing the neural network to learn the shape of the eclipse signal in contrast to the systematics. Figure 2 demonstrates the outcome of this scaling on a lightcurve dominated by a scattered light feature. The top panel shows the unscaled lightcurve, while the bottom panel shows the quantile-scaled lightcurve processed as input to the neural network. Although it is clearly more difficult for the human eye to distinguish eclipses in the scaled form, this method proved to be superior to the neural network for understanding subtle differences between eclipses and eclipse-shaped noise or systematics. This is not to say that we were able to avoid outcomes where the neural network classifies such features as eclipses entirely, but this method did substantially diminish the problem.

While quantile scaling underemphasizes large features, it also has the effect of overemphasizing small features. To our benefit, this helped in the identification of shallow eclipses.

However, we also found that our network will identify planet transits as well as small eclipse-like shapes in noise patterns, which became a substantial source of error (discussed further in Section 3.3). Thus, we effectively made the decision to trade large noise effects for small noise effects. We make no assertion that this trade was ideal, nor our method of classifying the eclipse shape versus identifying an EB directly.

3.2. Training Data Collection

A particular challenge of this effort was the collection and augmentation of the training data set. Generally, the performance of a classifier will track directly with the quantity of training data samples to an asymptotic limit (e.g., C. Sun et al. 2017). Additional difficulties arise for our particular application in that eclipses can be vastly different in appearance, thus requiring a substantial amount of training




Figure 2. Lightcurve of TIC 139079180 for Sector 15 (top panel) vs. the same lightcurve scaled with a quantile transform (bottom panel). Eclipses are highlighted in blue. The scattered light near the ends of the two segments dominates the lightcurve and also resembles an eclipse, making this type of feature difficult to overcome as a source of false positives. The quantile transform has the effect of making these events less prominent and emphasizing the actual eclipses. The inputs to the neural network are the quantile-scaled lightcurves. Although making classification perhaps more difficult to the human eye, the quantile transform represents the lightcurve to the neural network in a manner that allows it to successfully find and classify the eclipse feature.

data for a neural network to effectively generalize the features of an eclipse.

At the time of development of our neural network (TESS Year 2), we had tens of millions of lightcurves, but only a handful of manually selected EB lightcurves. To gather a sufficiently sized data set to effectively train our neural network by manually sorting through individual lightcurves would have been an intractable task. As such, we progressively augmented our data set by iteratively using a weakly trained neural network to find new training samples among the full data set of lightcurves. After each iteration of training and inference, we would select (i) lightcurves given a score near unity that clearly did not show an eclipse, or (ii) lightcurves given a score near zero that clearly showed an eclipse. The former represented false positives and the latter false negatives. We used these as properly labeled training samples in the next iteration of training, effectively filling gaps in the understanding of the neural network. With each of these

iterations, the neural network became progressively more capable. By the time we were satisfied with the performance of the neural network, we had built a training set of $\sim 40,000$ samples.

3.3. Model Performance

We emphasize that our neural network was not trained to find EBs. It has no concept of repeated features or periodicity. Rather, it was trained to find eclipses, or, more broadly, features resembling eclipses. We show how the neural network activates on the shape of the eclipse in the saliency map of the activation weights of the penultimate layer in Figure 3. Note that the neural network will emphasize a single eclipse in determining the output score of the lightcurve.

Knowing that the neural network would be providing candidates for manual review rather than directly populating a list of near-certain EBs, we wanted to allow for interesting results that would not fit the conventional shape of an eclipse,




Figure 3. Neural network saliency map (shades of red according to the activation magnitude) for a segment of the TIC 214716930 lightcurve in TESS Sector 12 (blue), demonstrating the activation of the neural network in the penultimate convolutional layer on the feature of the eclipse, made using KERAS-VIS (R. Kotikalapudi 2017).

e.g., a complex syzygy or lopsided eclipse. This process has allowed us to find multiple star systems with complex outer orbital eclipses, among other interesting phenomena, the body of work on which has established the effectiveness of our methods given that this neural network has contributed to many discoveries (e.g., V. B. Kostov et al. 2021a, 2021b, 2022b, 2023, 2024a, 2024b; B. P. Powell et al. 2021a, 2021b, 2022b, 2023, 2025; T. Borkovits et al. 2022; B. K. Capistrant et al. 2022; S. A. Rappaport et al. 2022, 2023, 2024; R. Jayaraman et al. 2024; T. Mitnyan et al. 2024a; K. Oláh et al. 2025).

Our results have been qualitative, with a manual review of outputs in multiple stages (described further in Section 4). As such, an assessment of performance of the neural network against a data set of known EBs is hardly direct. However, we provide such a comparison here to provide the reader with the context of our process as well as the contents of our catalog.

An evaluation of the model would be most complete with a section of the sky where we could consider all EBs within TESS’s limiting magnitude to be known. As such, the Kepler (W. J. Borucki et al. 2010) field provided an ideal testing ground, with the full data set having been thoroughly evaluated for the presence of EBs, resulting in the production of an EB catalog (A. Prša et al. 2011; R. W. Slawson et al. 2011; G. Matijevič et al. 2012; K. E. Conroy et al. 2014a, 2014b; D. M. LaCourse et al. 2015; M. Abdul-Masih et al. 2016; B. Kirk et al. 2016), hereafter referred to as the “Kepler EB catalog.” By comparing the number and characteristics of lightcurves identified in our catalog against those of the 2920 EBs of the Kepler EB catalog, we could make an estimate of the performance of our catalog.

We also considered the catalog of 4584 TESS EBs from short-cadence data in Sectors 1–26 (A. Prša et al. 2022), hereafter referred to as the “TESS EB catalog.” Although less comprehensive in terms of a full survey of a section of the sky, a direct comparison to TESS EBs rather than Kepler allows for

fewer independent sources of error such as different noise amplitudes or photometric capabilities.

We cross-referenced our catalog with the Kepler field boundaries,²² finding each object from our catalog that would have been observed by the original Kepler mission, which resulted in 9768 unique TIC IDs. We then reduced the Kepler EB catalog to those objects with $T_{\text{mag}} < 15$, as this was the limit of our lightcurve construction, resulting in 2458 EBs out of the original 2920. Our catalog contains 1371 of these 2458 EBs, or $\sim 55.8\%$. To compare our catalog to the TESS EB catalog, we cross-referenced our catalog with lists of the 2 minute targets for TESS Sectors 1–26,²³ from which the TESS EB catalog was derived. In total, there were 507,898 unique 2 minute targets in these sectors, 8910 of which were identified by our neural network. Comparing this sample directly to the TESS EB catalog revealed that our neural network found 3884 of the 4584 EBs therein, or $\sim 84.7\%$. In this comparison of true positives, it is clear that our neural network performed far better than the TESS EB catalog, which we assess to be likely due to the systematic differences between the TESS and Kepler data.

In Figure 4, we show a scatter plot of the morphology parameter versus T_{mag} of the Kepler EBs found (blue) and not found (red) by our neural network. The morphology parameter (described further in A. Prša et al. 2011) is a measure of the EB type, with values less than 0.5 corresponding to detached EBs, values in the range 0.5–0.7 corresponding to semidetached EBs, values in the range 0.7–0.8 corresponding to overcontact EBs, and values greater than 0.8 corresponding to ellipsoidal or unknown classifications. We exclude analysis where the morphology parameter was given a value of -1 , indicating the lightcurve was unclassifiable by the methods of A. Prša et al.

²² Available from MAST at https://archive.stsci.edu/missions/kepler/ffi_footprints/morc_2_ra_dec_4_seasons.txt.

²³ https://tess.mit.edu/public/target_lists/target_lists.html




Figure 4. (Left panel) Scatter plot of the morphology parameter vs. T_{mag} for the Kepler EBs found (blue) and not found (red) by our neural network. (Right panel) The same plot for the TESS EB catalog. For both panels, the top histogram shows the distribution over morphology, while the right histogram shows the distribution by T_{mag} . Our performance against the Kepler EBs shows a clear preference for the central morphology range as well as a decline as magnitude decreases. Note that the T_{mag} distribution for the TESS EBs is limited by their selection as 2 minute cadence targets, hence the apparent cutoff at $T_{\text{mag}} \approx 12$.




Figure 5. (Left panel) Scatter plot of the EB period vs. T_{mag} for the Kepler EBs found (blue) and not found (red) by our neural network. (Right panel) The same plot for the TESS EB catalog. For both panels, the top histogram shows the distribution over EB period, while the right histogram shows the distribution by T_{mag} . Again, our performance against the Kepler EBs shows a decline as magnitude decreases. As with Figure 4, note that the T_{mag} distribution for the TESS EBs is limited by their selection as 2 minute cadence targets, hence the apparent cutoff at $T_{\text{mag}} \approx 12$.

(2011). Furthermore, we make the same comparison using the EB period versus T_{mag} in Figure 5 and the EB period versus the morphology parameter in Figure 6. We can make two conclusions from these figures:

(i) The T_{mag} histograms in the left panels (Kepler comparisons) of Figures 4 and 5 show a clear decrease in performance with decreasing T_{mag} , while we do not see the same decrease in the right panels (TESS comparisons) of the same figures. Although there is a somewhat artificial lower limit on magnitude in the TESS 2 minute cadence targets due

to the selection bias for bright stars, we can still see a generally uniform trend of the fraction identified in the T_{mag} histograms of the TESS comparisons. However, these targets are still somewhat idealized in comparison to a full sample over a section of the sky. As such, we consider our neural network's performance against the TESS EB catalog 2 minute cadence targets to be an upper limit, while the performance against the Kepler EB catalog should be considered a lower limit.

(ii) For the Kepler (left) panels in Figures 4 and 6, the morphology histograms show a clear weakness of our neural




Figure 6. (Left panel) Scatter plot of the EB period vs. the morphology parameter for the Kepler EBs found (blue) and not found (red) by our neural network. (Right panel) The same plot for the TESS EB catalog. For both panels, the top histogram shows the distribution over EB period, while the right histogram shows the distribution by morphology parameter. Besides revealing the relationship between EB period and morphology as described by A. Prša et al. (2011), we can see a pattern of weakness of our neural network for EBs at the extremes of the period range for any given morphology range.

network for the extremes of the morphology range. Aside from showing the general relationship between EB period and morphology, Figure 6 demonstrates a clear trend of weakness for our neural network in identifying EBs at the extremes of the period range for any given morphology range in both the Kepler (left) and TESS (right) comparisons. That is, the red points seem to dominate the blue both to the far left and to the far right of the general trend line. We assess that both of these trends demonstrate a weakness in generalization of the neural network to less common types of eclipse patterns.

Having examined the nature of our true positives and false negatives, we turned to the much larger set of false positives. As previously discussed, 1371 of the 9768 unique TIC IDs from the Kepler field and 3884 of the 8910 in the TESS Sector 1–26 short-cadence targets found by our neural network were true positives, leaving the remaining 8397 ($\sim 86\%$) of the Kepler sample and 5026 ($\sim 56\%$) of the TESS sample as false positives. Naturally, the questions arise as to what these false positives are and why they are so numerous. To determine their nature, we manually examined a subset of bright false positives in the Kepler field with $T_{\text{mag}} < 10$, totaling 126 unique TIC IDs.

Of these lightcurves, we found that 16 ($\sim 13\%$) showed clear EBs. Since these were not in the Kepler EB catalog, we assumed that most of these are likely the result of blending. That is, the $21''$ pixels of TESS will frequently cause the lightcurves of bright stars to show in the lightcurves of dimmer close neighbors. To demonstrate this effect, we show the Kepler field with true positives, false positives, and false negatives in Figure 7, where the clustering of several large groups of false positives can be seen, likely as a result of blending. While this was only a problem in $\sim 13\%$ of our false positives with $T_{\text{mag}} < 10$, it can be reasonably expected that lightcurves of dimmer stars will show this type of contamination more frequently. We examined each of the 16 lightcurves

showing clear EBs and confirmed that contamination from nearby brighter stars was indeed the source of the signal. However, we assessed two of the 16 lightcurves showing clear eclipses as the true source of an EB, TIC 26542657 (KIC 12013550) and TIC 63454475 (KIC 10342012). We confirmed that neither of these targets are present in the Kepler EB catalog. Furthermore, we found that there was no Kepler data available for TIC 63454475, while the Kepler lightcurves for TIC 26542657 (first identified as an EB in the TESS EB catalog by A. Prša et al. 2022) indeed showed no eclipses, confirming that neither of the targets were missed accidentally in the creation of the Kepler EB catalog. Given that TIC 26542657 has Kepler lightcurves without eclipses, we assess that it must be a higher-order system, likely a triple, with so-called “disappearing eclipses.” We will discuss this particular system further (also confirming it as a triple) as well as identify other examples of this type of system in Section 5.4. Although it is beyond the scope of this effort to analyze these systems, we note briefly that the changes in binary inclination causing periods of eclipsing and noneclipsing behavior are driven by interaction with the outer body, and we refer the reader to T. Borkovits (2022) for a detailed discussion of the nature of this type of triple system, among others. Given we found one such system out of only 126 in our crossmatch with the Kepler EB catalog with $T_{\text{mag}} < 10$, we expect there to be several more such systems in our 8397 false positives from the Kepler field, which may merit an investigation in its own right.

Returning to our false positives from the $T_{\text{mag}} < 10$ sample, oscillations resembling eclipses comprised 16 ($\sim 13\%$) of the false positives, while 21 ($\sim 17\%$) showed scattered light systematics resembling eclipses, as in Figure 2. The latter were particularly prominent in TESS Sector 14, which overlapped with the Kepler field. The bulk of the false positives, 74 ($\sim 59\%$), contained noise patterns that resembled eclipses. We provide examples of these false positives in Figure 8.




Figure 7. As compared to the Kepler EB catalog, true positives (blue), false negatives (red), and false positives (green) identified by our neural network, shown here in the Kepler field. Local groupings of false positives are attributable to very bright EBs contaminating the lightcurves of adjacent stars or systematics such as scattered light, which closely resembled the shape of an eclipse in the overlapping TESS Sector 14.

Another type of scientifically valuable false positive found by our neural network was planet transits. We compared our catalog to the NASA Exoplanet Archive Kepler catalog and found an overlap of 1502 out of the 8397 TIC IDs. Of these, 125 were determined to be genuine exoplanet candidates. The overlap with the NASA Exoplanet Archive TESS exoplanet candidate catalog²⁴ is even more pronounced. Of the 7576 TESS exoplanet candidates, our catalog contains 2445 ($\sim 32\%$). Given how we scale the lightcurves (see Figure 2), the neural network is generally not able to discern a difference between an eclipse and a clean transit signal, as it was not trained to do so. As such, we expect that our catalog is also rich with planet candidates. Of particular interest is the fact that, since our neural network has no periodicity requirement,

we expect there to be many single transit events that evade discovery through periodicity-based analyses.

3.4. Limitations and Caveats of Our Results

We emphasize, again, that our neural network was trained to find features resembling eclipses, not EBs. Our results were qualitative and manually reviewed. As such, the comparisons provided in the previous section should not be considered a fully accurate measure of the performance of the neural network as much as context for the reader to understand our process and the contents of our catalog. Our method should be considered as a means of reducing an extremely large data set (hundreds of millions of lightcurves) to a much smaller, manageable data set with a high concentration of scientific value.

Our full catalog consists of 1,223,603 unique TIC IDs with lightcurves that our neural network gave a score of ≥ 0.9 . The distribution of these, in terms of TESS magnitude and ecliptic

²⁴ Both Kepler and TESS exoplanet catalogs are available at <https://exoplanetarchive.ipac.caltech.edu/index.html>.




Figure 8. Examples of several different types of false positives returned by our neural network shown by TESS Sector 14 lightcurves. The TIC 120693310 lightcurve (upper left) demonstrates scattered light resembling eclipses, a particularly difficult false positive to train against. The TIC 120426180 lightcurve (upper right) shows pulsations with local minima broad enough to be classified as an eclipse. The TIC 26656569 lightcurve (bottom left) shows a clear EB, but this lightcurve is contaminated by the nearby TIC 26656583 (KIC 11560447), which is in the Kepler EB catalog. The TIC 123447105 lightcurve (bottom right) shows a noise pattern with several features that could be mistaken for eclipses.




Figure 9. Left: number of TESS targets exhibiting eclipse-like features from Sectors 1–82 as a function of TESS magnitude. Right: corresponding R.A. vs. decl.

coordinates, is shown in Figure 9. Most of the targets are on the faint side, with a median TESS magnitude of ≈ 14 , and the majority are near the Galactic plane. It is from these candidates that we distilled the much smaller catalog of vetted EBs to be discussed in the remainder of this paper. This catalog could be employed by interested researchers with the caveats of the analysis in the preceding section. We use the Kepler EB catalog comparison as the lower end of our estimate and the TESS EB catalog comparison as the upper end of our estimate in the following summary of caveats and contents:

1. This is a catalog of lightcurves with eclipse-like features, not an EB catalog.
2. A total of 14%–44% of the catalog should be expected to be EBs.
3. It follows that 56%–86% of the catalog should be expected to not be EBs. These will consist of contaminated lightcurves, transiting exoplanets, dippers, systematics, and other sharp ellipsoidal features.
4. The catalog should be expected to contain 55%–84% of all EBs in the TESS lightcurves for Sectors 1–82 with

$T_{\text{mag}} < 15$. The fraction of completeness will decrease with brightness.

5. The full catalog is entirely unvetted and we offer no guarantee as to its contents. We have, however, found the neural network outputs to be scientifically valuable, so we offer it to the community in its entirety for their own purposes.

With these caveats in mind, we transition in the next section to explaining how we analyzed a subset of these, producing 10,001 genuine EBs for inclusion in a well-vetted catalog.

4. Vetting and Validation of the EB Candidates

Contamination from nearby (in terms of sky projection) EBs can result in a not-insignificant contribution of light to the aperture used to extract the target’s lightcurve—and thus mimic eclipse-like signals that seem to come from the target star. This is a common occurrence in TESS observations, where it is not unusual to see one or more field stars within 2–3 pixels of the target star (where each pixel is $\approx 21''$), often even falling within




Figure 10. Example 5×5 pixel sector-averaged difference images used for photocenter measurements of TESS EB candidates. The red star symbols represent the pixel position of the target star, the open black circles represent the sector-averaged photocenter, and the small red symbols represent nearby field stars that are bright enough to produce the detected eclipses as contamination to the target star. The first three columns from the left (labels (A), (B), and (C)) show difference images that are reasonably well suited for photocenter measurements. The photocenters in the first column (A) indicate that the detected eclipses are on-target, whereas the second (B) and third (C) columns show false positives due to significant photocenter offset. Column (B) highlights the case for a well separated target and contaminator, while the two are within the same pixel in column (C)—and also reside in crowded fields. Column (D) shows difference images that are dominated by systematic effects, thus making them inadequate for reliable photocenter measurements.

the same pixel, and thus adding their often significant signal to that of the target of interest (e.g., C. X. Huang et al. 2020a, 2020b; V. B. Kostov et al. 2022b, 2024a; M. Kunimoto et al. 2024; and references therein). Thus detection of eclipse-like features from a particular star in TESS does not immediately tell us their origin and additional investigations are required before an EB candidate is verified. In the absence of radial velocity measurements to confirm or rule out a potential EB (or planet) candidate, one can capitalize on the information-rich content of the available photometric data.

4.1. Photocenter Vetting

A particularly powerful method to constrain the pixel position of the source of detected eclipses (or transits) is the photocenter-based analysis that is routinely used in surveys aimed at finding transiting exoplanets (e.g., J. L. Coughlin et al. 2014; S. E. Thompson et al. 2015, 2018; J. D. Twicken et al. 2018; V. B. Kostov et al. 2019; D. J. Armstrong et al. 2021; H. Valizadegan et al. 2025; and references therein). Briefly, the method uses the target pixel files to measure the in-eclipse center-of-light (“photocenter”) pixel position for each eclipse detected in the difference image,²⁵ and compare it to the out-of-eclipse photocenter and/or to the catalog pixel position of the target. If there is no statistical difference between these and the in-eclipse photocenter, the eclipses are considered to be “on-target”; otherwise the eclipses are “off-target,” likely coming from a nearby field EB, and the candidate is marked as a false positive.

Obtaining robust photocenter measurements depends on multiple factors such as the signal-to-noise ratio (SNR) of the target pixel files, the depth of the detected eclipses, the presence of nearby comparably bright field stars (worst-case scenario: ones that are much brighter and variable on timescales comparable to the duration of the eclipses), etc. In theory, the difference images used to measure the per-eclipse photocenters resemble a well-defined, bright pixelated spot superimposed on an otherwise dark background (Figure 10, first three columns from the left). In practice, the difference images are often distorted due to various astrophysical, systematic, or instrumental effects (Figure 10, rightmost column), making the corresponding photocenter measurements unreliable (e.g., L. Caciapuoti et al. 2022; V. B. Kostov et al. 2022b, 2024a; C. Magliano et al. 2023).

Overall, based on our experience with TESS data—and depending on the peculiarities of the specific target—measurements of genuine photocenter offsets of $\gtrsim 0.2$ – 0.3 pixels (i.e., $\gtrsim 4''$ – $6''$) are often trustworthy (V. B. Kostov et al. 2022b, 2024a; C. Magliano et al. 2023). However, considering measured offsets of $\lesssim 0.1$ – 0.2 pixels as significant can be extremely challenging or potentially even impossible. Thus, throughout this work we adopt a photocenter offset threshold of 0.2 pixels ($\approx 4''$) such that cases below are considered as likely “on-target” and those above “off-target.”

Ideally, genuine photocenter offsets 3–5 times larger than this threshold (i.e., $\sim 12''$ – $21''$) should be relatively easy to measure (and accept as reliable). Thus, to account for potential false positives due to known EBs, we first evaluated whether the TESS EB candidates our neural network identified are within a $21''$ (1 TESS pixel) sky-projected separation of EBs listed in various catalogs. In particular, we queried EB catalogs from ASAS-SN

²⁵ The difference image is created by subtracting the out-of-eclipse target pixel data from the in-eclipse pixel data, both covering the measured duration of the eclipse.




Figure 11. Distributions of the TESS magnitudes, preliminary measured periods, and respective photometer offsets for the 60K sample. The lower right panel is a zoomed-in version of the lower left panel, highlighting the distribution of measured photometers smaller than 2 pixels. See text for details.

(~150,000 EBs), ATLAS (~30,000 EBs), Gaia (~2,200,000 EBs), OGLE (~430,000 EBs), Simbad (~2,400,000 EBs, ~200,000 spectroscopic binaries), TESS (~50,000 EBs, ~10,000 planet candidates from ExoFOP-TESS²⁶), the International Variable Star Index (VSX; ~900,000 EBs), and WISE (~50,000 EBs), taking into account the respective overlaps between the different data sets. Unsurprisingly, given the pixel size of TESS and the corresponding crowding, about one-quarter of our ~1.2 million candidates fulfill the above criteria (see Table 1). We note that this consideration does not immediately rule these out as false positives, but it marks them as likely suspects. Conversely, those that are not within 1 pixel of known EBs (~900,000 TIC IDs) can potentially be confirmed as bona fide new EBs through careful photometer analysis. For the benefit of the community, we provide these targets as an auxiliary data set (see Table 2).

To investigate this matter further, we conducted a deep dive into a subset of ~60,000 (hereafter 60K) targets, representing ~5% of the likely known and potentially new EB candidates from our preliminary list. The targets were randomly selected and evenly split on either side of the 21" demarcation line, and are representative of the TESS magnitude and sky position distributions shown in Figure 9. These 60K targets were subjected to comprehensive ephemeris determination and photometer measurements, and analyzed in depth via the two-step process outlined below.

4.2. In-depth Analysis of 60K Targets

During the first step, we developed and utilized an automated pipeline to calculate ephemerides and measure photometer offsets. Specifically, we applied the box least-squares algorithm

²⁶ <https://exofof.ipac.caltech.edu/teess/>

Table 1
Likely Overlap between TESS Targets Exhibiting Eclipse-like Features
Detected by Our Neural Network and Known EBs

Source	Number of Targets
This work	1,223,603
ASAS-SN	119,126
ATLAS	21,244
Gaia	259,631
OGLE	6841
Simbad	268,669
Simbad (spectroscopic binaries)	14,076
TESS (EBs or planet candidates)	62,681
VSX	136,190
WISE	41,085
ZTF	18,336
Total overlap	343,014

Notes. “Overlap” is defined here as a sky-projected separation of less than 1 TESS pixel (~21"). Duplicates are removed from the total number of overlaps.

(BLS; G. Kovács et al. 2002) to the available ELEANOR FFI lightcurves to measure periods and conjunction times, limiting the minimum/maximum period searched for to 0.5/40 days, respectively. The BLS results were further improved by fitting a generalized Gaussian model to each detected eclipse adopting the methodology of V. B. Kostov et al. (2022b), and testing for period deviations from the linear ephemeris. The latter helps take into account potential eclipse timing variations (ETVs) that may decrease the precision of the BLS measurements, and also provides robust measurements for the eclipse depths and durations. Next, we used the refined ephemerides and eclipse durations to construct the appropriate difference images for each EB candidate, following the prescription of V. B. Kostov et al.




Figure 12. An example issue associated with the automated ephemeris measurements of the TESS FFI ELEANOR lightcurve and difference image for TIC 33521833. The upper panel shows one sector of TESS data with the automatically detected periodic signal highlighted by the green vertical bands. Lower left panel: corresponding phase-folded data from all available sectors. Lower right panel: corresponding difference image and photometer measurements. Here, the lightcurve is dominated by stellar variability and the automatically measured period is incorrect.

Table 2

Identifying Information for $\sim 900,000$ Unvetted and Unvalidated Targets That Match Our Selection Criteria

TIC ID	R.A. (deg)	Decl. (deg)	T_{mag}
1051	218.815978	-28.267080	14.77
4482	218.858361	-25.722840	13.42
8639	219.017912	-27.561259	14.75
17084	219.309771	-25.415941	14.68
17361	219.336321	-24.958481	11.34

Note. Selection criteria: (i) the neural network identified eclipse-like events with a score greater than 0.9, and (ii) no EBs from the sources listed in Table 1 are within $\sim 21''$.

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

(2019). Finally, we obtained photometer measurements by fitting to each difference image the TESS pixel response function and a Gaussian point-spread function, and adopting the average of the two as the corresponding photometer of the image.

Preliminary results from the first step are highlighted in Figure 11, showing the distributions of the 60K sample in terms of TESS magnitude, measured period, and photometer offset. At this stage, the relevant measurements have not yet undergone the rigorous vetting and validation analysis required for promoting a target as a genuine EB, and are thus likely affected by various systematics. For example, the period distribution of the new EBs seen in Figure 11 shows a local maximum near 14 days. This is close to half the duration of a TESS sector, which makes the potential periods suspicious.

The second step of the process addresses the issue of potentially incorrect ephemeris measurements and, consequently, incorrect photometer measurements produced by the automated pipeline outlined above. This can occur when the period search is misled by, for example, strong systematic effects such as prominent out-of-eclipse lightcurve variations (due to, e.g., starspots) that can dominate or even completely overwhelm the eclipse signal (see Figures 12 and 13, upper panels). Additionally, EBs where the primary and secondary eclipses are similar in depth, duration, and shape can lead to situations where the automatically measured period is an integer ratio of the true period (see Figure 13, lower panels). While we tried to minimize the impact of such issues on the ephemeris and photometer pipeline as much as possible, it was challenging to account for all possible complications without negatively affecting the signals of interest. For example, we only used data points with good quality flags as provided by ELEANOR; removed those that were near known issues, such as momentum dumps (TESS Instrument Handbook²⁷), or that we identified as potentially suspect sections of the lightcurve; and, where appropriate, utilized low-order polynomial detrending.

To account for these and other challenges, we manually inspected the products of the automated pipeline for each of the 60K targets as outlined below.

4.3. Citizen Science

Analyzing vast amounts of data using automated methods remains a highly nontrivial process as various sources of “noise”—be it astrophysical, instrumental, or systematic—can introduce subtleties that are challenging to automatically

²⁷ https://archive.stsci.edu/files/live/sites/mast/files/home/missions-and-data/active-missions/teess/_documents/TESS_Instrument_Handbook_v0.1.pdf




Figure 13. Same as Figure 12 but for TIC 823000. Here, the primary and secondary eclipses have similar depths (130 parts per thousand (ppt) and 120 ppt, respectively) and the automatically measured period is off by a factor of 2.

account for. Additionally, while autonomous methods excel at, for example, finding a periodic signal in time series such as the transits of an exoplanet (e.g., using BLS), they often lack the insight to discover unique and unexpected features, or various unusual and uncommon astrophysical objects or phenomena. For example, all known transiting circumbinary planets have been identified by visual inspection of Kepler and TESS lightcurves (V. B. Kostov 2023). Notably, while human intervention can help fill in the voids left by algorithms, given the complexity and scope of astronomical data sets it is quite challenging for a traditional research group to manually inspect hundreds of thousands of targets.

Thankfully, to the rescue come volunteers from all walks of life that boost the capacity of bandwidth-limited professional astronomers manyfold and help tackle the ever-increasing volume of publicly available astronomical data. This so-called citizen science approach is not a new concept—professional and amateur astronomers have a fantastic and strongly intertwined history. One famous example is the “Harvard Computers” project, where some of the participants started with no formal astronomy training yet helped revolutionize astronomy and became some of the most successful professional astronomers (e.g., S. Nelson 2008). Among the more recent examples, many transiting planets with orbital periods longer than 1 yr have been discovered by citizen scientists (e.g., J. Wang et al. 2015), and hundreds of eclipsing triple- and quadruple-star systems have been spotted by eagle-eyed volunteers (e.g., T. Borkovits 2022 and references therein). Citizen scientists have been responsible for many “firsts,” e.g., (i) the unusual Boyajian’s Star (T. S. Boyajian et al. 2016), (ii) an exocomet transiting its host star (S. Rappaport et al. 2018), and (iii) a newly discovered class of objects called “tidally tilted pulsators” (e.g., G. Handler et al. 2020). Time and again, volunteers have demonstrated they can extract interesting signals from noise in numerous cases.

The state of citizen science is strong, with multiple projects tackling various astronomical data sets and making important scientific contributions on a regular basis thanks to, e.g., Planet Hunters and Planet Hunters TESS,²⁸ Exoplanet Explorers,²⁹ Citizen ASAS-SN,³⁰ SuperWASP Variable Stars,³¹ Planet Patrol,³² Eclipsing Binary Patrol (EBP),³³ Exoplanet Watch,³⁴ UNITE: Unistellar Network Investigating TESS Exoplanets,³⁵ and the Visual Survey Group (VSG; M. H. K. Kristiansen et al. 2022). The power and dedication of citizen scientists is truly incredible—for example, members of the VSG have visually inspected tens of millions of lightcurves from Kepler and TESS (M. H. K. Kristiansen et al. 2022), and helped make important new discoveries in multiple branches of astrophysics.

4.3.1. Exogram

The manual inspection of the 60K targets proceeded as follows. First, we adapted our online vetting portal Exogram³⁶ for rapid visual scrutiny of a randomly drawn subset of about 10,000 targets out of the 60K set by our core science team composed of professional astronomers and highly knowledgeable citizen scientist “superusers.” The team formed during the

²⁸ <https://blog.planethunters.org/2010/12/16/planet-hunters-introduction/>, D. A. Fischer et al. (2012); <https://www.zooniverse.org/projects/nora-dot-eisner/planet-hunters-tesse>.

²⁹ <https://exoplanetexplorers.org>

³⁰ <https://www.zooniverse.org/projects/tharinduj/citizen-asas-sn>

³¹ <https://www.zooniverse.org/projects/ajnorton/superwasp-variable-stars>

³² <https://www.zooniverse.org/projects/marckuchner/planet-patrol/>

³³ <https://www.zooniverse.org/projects/vbkostov/eclipsing-binary-patrol>

³⁴ <https://exoplanets.nasa.gov/exoplanet-watch/about-exoplanet-watch/overview/>

³⁵ <https://science.unistellar.com/exoplanets/unite/>

³⁶ <https://exogram.vercel.app/>, developed by one of the citizen scientists on our team (R. Salik).




Figure 14. An example screenshot of the workflow and interface of Exogram EB vetting, enabling rapid image classification.

Planet Patrol project, and we have been working together ever since. The custom interface breaks down the vetting process into three main questions: “Is this an EB?,” “Is the measured period correct?,” and a space for additional comments (both predefined and free text). An example screenshot from the Exogram EB vetting portal is shown in Figure 14.

Briefly, the user first evaluates the data for clear evidence of eclipse-like features, paying close attention to the two panels on the right, and to the lower left panel. Next, Exogram automatically proceeds to the second question, where the user scrutinizes the phase-folded plot (middle left panel) and decides whether the measured period is correct. Finally, the portal takes the user to the last question, which provides the opportunity to mark the target as particularly noteworthy. Several predefined options are provided, corresponding to the most commonly observed characteristics.

Importantly, Exogram is designed to enable fast image classification through the use of keyboard shortcuts. We found this to be a critical advantage as it allows an expert vetter to classify images with a typical “cruising speed” on the order of seconds, especially when the data clearly indicates a typical EB system.³⁷ Naturally, more interesting cases such as those exhibiting additional eclipse-like features take longer to inspect, as do targets dominated by systematics such as momentum dumps, but even for these the access to keyboard shortcuts significantly decreases the response time.

The portal also provides links to external tools such as the Fast Lightcurve Inspector (FLI³⁸) and LATTE (N. Eisner 2022) that enable deeper investigation of potentially interesting or particularly challenging targets. Both tools allow researchers to interactively examine the entire lightcurve, making it easier, for example, to distinguish between genuine eclipses and momentum dumps or other sources of interference. It is worth

noting that FLI was created by one of our superusers (J. S. de Lambilly) and is designed as a free, online, user-friendly, interactive tool for visual inspection of TESS data, including BLS analysis and phase folding. FLI uses the Lightcurve package (Lightcurve Collaboration et al. 2018) to query MAST for all available GSFC-ELEANOR-LITE, QLP, SPOC, and TESS-SPOC lightcurves, along with corresponding diagnostics (e.g., background flux and centroid measurements), and presents them in a Bokeh/Plotly environment.

Exogram was developed using SvelteKit, a modern web application framework. The database and authentication are handled by Google’s Firebase platform, and the website itself is hosted by Vercel. We store the lightcurve images on Google Drive. Behind the scenes, the Exogram server tallies the number of responses, labels targets as fully classified if three users have already commented, and removes them from the pool of images shown. Additionally, the Exogram platform integrates social media-like features to encourage collaboration between users. For example, users can “star” a target to save it for later inspection. Starred EBs are public, and users can see which targets were saved by others. This makes it easy for vetters to find targets that others deemed interesting or unusual. Users can also share targets with each other, and the integrated notification feature alerts users when something has been shared with them. Finally, the platform also shows a vetting leaderboard to encourage friendly competition among the users.

4.3.2. EBP

To further capitalize on the power of citizen science, and inspired by the success of the Planet Patrol project (V. B. Kostov et al. 2022a), we proceeded with the investigation of all 60K targets by developing the EBP project.³⁹ EBP is hosted on Zooniverse and provides a streamlined, interactive, and user-

³⁷ For the easiest cases, the vetting speed is practically limited by the reaction time for pressing the relevant shortcut.

³⁸ <https://fast-lightcurve-inspector.osc-fr1.scalingo.io>

³⁹ The targets inspected through Exogram are included in EBP as an additional layer of scrutiny and validation.




Figure 15. An example screenshot highlighting the first question of the EBP workflow. The upper panel in the figure shows one sector of TESS data, highlighting the detected eclipses in green, and listing the measured period. The lower left panel shows the corresponding phase-folded lightcurve, and the lower right panel shows the difference image used for the photocenter measurements. The difference image also shows the pixel position of the target (red star) and the sector-averaged measured photocenter (open black circle). The user decides whether the target shows clearly visible periodic dips (indicating an EB), answers “yes” or “no,” and is then taken to the second question.

friendly platform to visually inspect a summary of the results produced by the automated pipeline described above. EBP was launched on 2024 September 3, and was completed on 2025 March 26, during which period ~ 1800 participants produced $\sim 320,000$ classifications.

The EBP workflow consists of four questions aimed at evaluating whether the target is indeed an EB candidate, whether the calculated period is correct, and whether the photocenter measurements are reliable. An example screenshot of the classification scheme is shown in Figure 15. The volunteers inspect the original and phase-folded lightcurves, and decide whether they see periodic eclipse-like signals, check if the period is correct, and scrutinize the lightcurve for secondary eclipses. Additionally, they evaluate the quality of the difference image and classify it as either appropriate for photocenter measurements—i.e., the image shows a well-defined bright spot on an otherwise dark background—or otherwise.

The EBP portal provides extensive background information on the science of EBs and their astrophysical importance, a comprehensive tutorial with a step-by-step demonstration of the workflow, a field guide presenting relevant examples, edge cases, etc. The portal also includes guidelines on how to interpret and classify the images, as well as an active talkboard where volunteers can discuss targets of interest and ask for help from the science team. Each image presented to the volunteers also contains additional auxiliary information enabling more detailed investigation of the inspected target, particularly with the help of FLI. We note that classifications on EBP are not strictly blind. Volunteers could freely look up outside information, based on the provided TIC ID, which could potentially affect their evaluation. Finally, volunteers interested in contributing to the vetting process beyond the Zooniverse project are invited to join the science team.

For completeness, we would like to briefly share the experience gained and lessons learned from EBP. First and foremost, frequent interactions between the science team and

volunteers on the talkboards, especially during the initial stages, were critical for the success of the project. These interactions ensured timely resolution of technical issues, addressed vetting and scientific questions, and enabled live updates aimed at improving the overall workflow. For example, prompted by feedback from volunteers we quickly refined the FAQ and tutorial by adding new examples, clarifying existing instructions, etc. Finally, consistent communication, including social media posts highlighting interesting targets and celebrating milestones, helped retain user engagement throughout the duration of the project, averaging about 1000 classifications per day, even months after launch.

The workflow of EBP is designed such that each target is considered as fully classified when at least five different volunteers have inspected the corresponding image and answered the provided questions. It is worth pointing out that at the launch of the project, the image “retirement” limit was set to nine. However, that proved to be too high as the rate of completed classifications was rather slow. Thus, in order to speed up the vetting and complete the 60K sample in a timely manner, three weeks after the project was launched we reduced the limit to seven, and shortly after down to five.

To adopt an aggregate response to each question, we tested three options: (i) a simple majority, i.e., at least three out of five volunteers select the same answer; (ii) at least four out of five; and (iii) five out of five. To evaluate the reliability of these aggregates, we checked the corresponding classifications for a random sample of 1000 targets where the measured period was classified as correct. Overall, we found that about 75%, 85%, and 90% of the responses are correct for options (i), (ii), and (iii), respectively. In order to increase the fidelity and maximize the reliability of the classifications, members of the science team performed complementary visual inspection of all potentially new EB candidates for which at least three out of five volunteers indicate that the period is correct.

Altogether, 7936 targets passed the following vetting and validation tests: (1) the ELEANOR lightcurve shows clear




Figure 16. Same as Figure 12 but for TIC 187172446. Here, the lightcurve shows a clear EB signal, the measured period is correct, the difference image is adequate for reliable photometer analysis, and the measured photocenter offset is ~ 0.14 pixels. However, there is a nearby field star, TIC 510123334, that is bright enough to be the source of the eclipses and too close to the target (projected separation of about $0''.71$) for the photometer measurements to pinpoint the source of the detected eclipses. Such targets are not included in our catalog.




Figure 17. The distributions of the TESS magnitudes, orbital periods, photocenter offsets, and number of sectors observed for the 7936 new EBs in our catalog. A darker shade represents a higher number of targets. Most of the EBs are faint, have short orbital periods, have detected eclipses that originate within $\sim 1''$ – $2''$ of the respective target star, and have been observed by TESS in at least three sectors.

eclipses; (2) the measured period is correct; (3) the difference images used for photometer analysis are of sufficiently high quality for reliable measurements; (4) the measured photocenter offsets are smaller than 0.2 pixels; and (5) no field stars from the Gaia catalog and TIC are within 0.2 pixels of the target, and bright enough to produce the detected eclipses as contamination. An example of a target that passes the first four tests but fails the last is TIC 187172446, shown in Figure 16. Here, there is a nearby field star, TIC 510123334, that is about 1 TESS magnitude fainter and at a projected separation of $0''.17$. Thus, while the measured photocenter offset is about 0.14 pixels, it is impossible to tell from the TESS data which of these two stars is producing the detected eclipses.

5. A Catalog of Uniformly Vetted and Validated EBs from TESS FFI Data

The final product of the process outlined above is a uniformly vetted and validated catalog of new EB candidates identified in TESS FFI data. The catalog contains 7936 targets with verified ephemerides, eclipse depths and durations, and, where applicable, phases of secondary eclipses. Table 3 shows the content of the catalog, which also includes the TIC ID of the target, sky position, TESS magnitude, number of sectors observed, Gaia astrometric information, relevant comments, etc. We note that 29 of the 7936 targets are listed in Gaia as single-lined spectroscopic binaries. In addition, we provide updated ephemerides for 2065 known EBs for which the period listed in one or more catalogs is incorrect. Most of our

Table 3
Parameters of the 7936 New EBs

TIC ID	R.A. (deg)	Decl. (deg)	T_{mag}	Period (days)	$T_{0,\text{prim}}$ (BJD ^a)	Prim. Depth (ppt)	Prim. Duration (hr)	$T_{0,\text{sec}}$ (BJD ^a)	Sec. Phase	Sec. Depth (ppt)	Sec. Duration (hr)	Sector	RUWE	AEN (μas)	AENS	T_{eff} (K)	Comments
8636	219.005555	-27.569441	12.36	3.886554	1603.2923	121	3.0	2355.3404	0.5	45	3.0	2	1.78	0.17	39.28	4965	...
672717	73.486075	-26.691694	14.44	7.18193	1448.2325	121	2.6	1445.1831	0.58	113	2.6	2	8.61	2.71	912.95	...	...
737910	74.543779	-28.491438	13.32	2.297452	1447.8866	31	3.6	1453.6230	0.5	30	3.5	2	1.06	0.07	1.08	4949	...
823000	127.043916	-16.737116	12.05	7.530827	2979.1752	132	3.5	2967.8872	0.5	121	4.6	4	0.99	0.00	0.00	5991	...
890432	127.043916	-13.796373	14.63	14.201932	2235.1764	143	4.5	2980.5780	0.49	58	4.9	5	1.12	0.09	8.63	4841	...
891636	127.057470	-12.512750	13.61	2.023998	1496.7517	99	2.1	2973.2540	0.5	38	2.7	5	39.40	4.10	11,889.59	...	...
1102444	155.353410	27.595890	12.8	4.04989	2612.6915	16	4.9	...	...	...	...	2	1.01	0.00	0.00	4087	...
1124666	36.948680	-7.826921	13.78	3.524638	2148.3515	139	3.5	...	...	...	...	2	...	6.11	9190.83	...	...
1195217	136.100062	-11.945334	13.29	1.788685	1541.2277	16	1.8	...	...	...	...	4	1.02	0.00	0.00	4156	...
1195846	71.095974	-32.854422	14.69	3.141602	2145.8888	33	2.5	2191.4521	0.5	8	2.7	3	1.02	0.00	0.00	4518	...
1196032	71.069268	-33.584354	14.64	15.15001	1444.4083	173	4.8	1449.5224	0.34	124	5.0	3	0.96	0.00	0.00	5516	...
1199301	71.333957	-35.113030	13.08	5.26431	2188.7857	62	6.0	1443.8504	0.49	3	6.1	3	0.98	0.00	0.00	6086	...
1309535	72.858580	-32.250238	12.12	13.531491	2187.5640	192	3.9	1452.1739	0.65	10	4.2	2	0.89	0.04	1.86	5987	...
1471956	146.911387	-15.546746	13.03	9.007584	2995.8203	135	3.3	2991.3347	0.5	81	3.3	4	18.99	2.10	2515.51	4326	...
1503836	131.959336	-24.266628	13.07	1.021917	1527.7987	21	1.6	2993.7410	0.5	12	1.7	7	1.92	0.21	34.13	...	...
1541478	169.541503	-14.092368	13.52	14.665098	3026.3663	417	3.1	3033.0097	0.45	359	3.0	4	1.19	0.05	0.76	3729	...
1616408	132.133251	-24.902464	14.71	3.115479	1524.3949	28	3.9	...	...	...	...	6	1.10	0.06	0.87	5861	...
1755837	78.624341	32.548174	14.28	4.28833	2502.8285	128	5.0	3645.6576	0.5	21	3.8	8	1.01	0.00	0.00	5368	...
1942820	78.947929	34.186537	14.34	3.381414	1826.4456	141	2.8	...	...	...	...	8	4.14	0.83	179.45	5780	...
2099994	79.341909	30.521062	10.91	3.827636	2496.3074	8	4.0	2521.1347	0.49	1	3.9	6	8.15	1.26	3159.49	5923	...
2158899	79.451105	30.893429	14.69	1.805987	2497.1628	20	3.5	...	...	...	...	8	1.06	0.04	0.31	6622	...
2438442	80.004098	35.227821	12.81	3.147522	1826.3919	133	4.0	3647.2077	0.49	55	3.8	7	18.74	2.18	3404.77	5972	...
2508333	80.125447	32.402184	13.85	2.183682	1840.1360	60	4.7	...	...	...	...	8	1.51	0.13	9.05	7104	...
2509102	79.993792	32.023402	11.78	2.486683	1825.6144	20	5.1	...	...	...	...	8	1.00	0.07	6.26	7141	...
2601042	80.292516	35.601431	13.57	3.476145	2493.7629	4	2.8	2481.5927	0.5	1	3.1	6	1.13	0.09	3.01	7681	...
2678574	80.370472	32.400301	13.59	1.329262	2475.2094	14	2.7	...	...	...	...	8	0.95	0.00	0.00	6825	...
2679754	80.290739	31.819113	12.32	3.184818	2508.1125	9	3.4	...	...	...	...	8	0.96	0.04	2.04	7065	...
2680580	80.337560	31.421225	14.28	6.874418	2492.9904	74	5.8	1829.4379	0.48	46	6.3	8	1.08	0.00	0.00	6650	...
2762701	356.676090	-16.842597	14.08	1.216169	1366.2573	179	1.3	1359.5735	0.5	69	1.4	3	1.05	0.08	0.75	...	pETVs
2840082	243.496472	-41.192966	11.49	2.887738	3073.6308	42	9.0	3092.4026	0.5	7	9.7	3	3.66	0.41	211.15	6456	...

Note. In the Comments column, pETVs = potential ETVs.

^a BJD-2457000.

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

Table 4
Comparison between the Correct Period Measured from TESS and the Periods from Gaia, ASAS-SN, ATLAS, VSX, and WISE for 2065 Known EBs

TIC	R.A. (deg)	Decl. (deg)	TESS Period (days)	TESS T_0 (BJD)	Gaia Period (days)	Gaia vs. TESS	ASAS-SN Period (days)	ASAS-SN vs. TESS	ATLAS Period (days)	ATLAS vs. TESS	VSX Period (days)	VSX vs. TESS	WISE Period (days)	WISE vs. TESS
761795257	111.7000	12.5605	2.5887	1494.1521	3.2064	1.33	...	...	5.1740	2	...	...	...	...
252351823	74.9920	55.4948	4.0213	1820.4185	7.4545	0	...	...	...	...	4.0088	1	...	...
436564213	70.6023	12.9477	3.2533	1444.2706	3.2431	1	3.2432	1	...	...	3.2431	1	1.6215	2
386250632	139.1054	−58.3239	1.5174	1546.1097	1.5174	1	9.1040	6	...	...	1.5172	1	0.7587	2
369995729	20.8105	59.8022	2.6458	1792.0862	2.0944	1.33	0.2571	0	...	...	0.2571	0	...	...
68543179	103.6037	32.4242	4.1924	1845.8610	4.6823	0	4.1927	1	...	...	4.1925	1	...	...
311651226	260.8442	−77.3195	1.0494	1629.9139	1.0494	1	1.0494	1	...	...	1.0524	1	0.5247	2
427654873	349.4827	70.1097	4.2869	1764.9825	8.5734	2	...	...	...	...	4.2871	1	2.1434	2
410498300	26.0089	48.0453	2.2905	1792.8908	2.2905	1	2.2903	1	...	...	2.2904	1	1.1453	2
123135027	118.6583	−5.2361	2.1691	1494.1874	0.2870	0	2.1689	1	...	...	2.1689	1	...	...

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




Figure 18. Example fits to all phase-folded primary eclipses of TIC 470715046 for three models: trapezoid (red), Gaussian (green), and generalized Gaussian (cyan). The red and green curves are vertically offset for clarity. As seen from the figure, the generalized Gaussian model provides a much better fit to the data than either of the other models.




Figure 19. Distributions of the median primary and, where present, secondary eclipse depths (left panel, in ppt) and durations (right panel, in hours) for the 7936 new EBs presented here.




Figure 20. Left panel: distribution of the detected secondary eclipses as a function of the orbital phase. Middle and right panels: the two EBs with the most extreme secondary phases, TIC 149673382 with phase ≈ 0.87 (middle) and TIC 337097515 with phase ≈ 0.11 (right panel), both exhibiting a heartbeat “bump.”

corrections are with respect to the Gaia EB catalog—1233 out of 1889 targets, followed by 312 out of 986 ASAS-SN EBs, and 308 out of 1015 VSX EBs.

The distributions of the TESS magnitudes, orbital periods, photocenter offsets, and number of sectors observed for the 7936 new EBs are shown in Figure 17. The period distribution has mean and median values of ≈ 4.5 and ≈ 3.5 days, respectively, and a 95th percentile of ≈ 11.6 days. This is comparable to the Kepler EB catalog, where the median period is also about 3.5 days (A. Prša et al. 2011; R. W. Slawson et al. 2011; K. E. Conroy et al. 2014a; W. F. Welsh &

J. A. Orosz 2018). The shortest period in our catalog is ≈ 0.65 days, while the longest is ≈ 40 days. As seen from the figure, the measured photocenter offsets are remarkably small, with mean/median/95th percentile values of 0.053/0.05/0.12 pixels, respectively, confirming that the detected eclipses originate from within $\sim 1''$ – $2''$ of the target stars. Additionally, granted that most of the EBs are on the fainter end (i.e., mean/median/95th percentile of 13.6/13.8/14.9 mag) there is no strong correlation between a target’s brightness and the magnitude of the corresponding photocenter offset, highlighting the excellent quality of the TESS FFIs even for the




Figure 21. Phase-folded lightcurves for two targets with primary eclipse depths close to unity: TIC 42066695 (sector-averaged depth of ≈ 848 ppt, left panel) and TIC 446208053 (sector-averaged depth of ≈ 804 ppt, right panel).




Figure 22. TESS FFI ELEANOR lightcurve of TIC 5232381 for Sector 9 (left) and Sector 62 (right). The vertical span is the same for both panels. The eclipse depths are different between the two sectors due to systematic effects caused by contamination from TIC 5232374.




Figure 23. Similar to Figure 22 but for TIC 77392704. The vertical span is the same for all three panels. Here, the eclipses are loud and clear in Sectors 12 and 39, but barely present in Sector 65.

fainter stars. Finally, TESS observed the majority of the EBs three times or more, such that $\approx 88\%/32\%/13\%$ of the targets were covered in at least three/six/nine sectors, respectively.

5.1. Depth, Duration, and Secondary Eclipses

To measure the eclipse times, depths, and durations, we adopted the methodology of V. B. Kostov et al. (2022b) and,

for each sector, fit each eclipse with its generalized Gaussian model:

$$F(t) = A - Be^{-\left(\frac{t-t_0}{\omega}\right)^\beta} + C(t - t_0). \quad (1)$$

For illustrative purposes, Figure 18 shows the model fit to all phase-folded primary eclipses of TIC 470715046, as well as a Gaussian and a trapezoid fit for comparison. As seen from




Figure 24. TESS FFI ELEANOR lightcurve of TIC 173706211 for all available sectors at the time of writing. Sector 84 is completely dominated by systematics. As a result, the ephemeris measured from Sectors 16 and 17 is not significantly improved with the addition of the latest data.

the figure, the generalized Gaussian model provides an excellent fit to the data—certainly better than both the trapezoid and the narrower Gaussian model—and we used it for measuring the eclipse depths and durations.

Table 3 provides the median depths and durations for the new EBs presented here. The corresponding distributions are shown in Figure 19. The primary depth distribution has mean/median/95th percentile values of 91/62/276 ppt, respectively; the mean/median/95th percentile values for the primary duration distribution are 4.2/3.9/7.4 hr, respectively. Roughly half of the targets exhibit secondary eclipses. As highlighted in Figure 20, most of these occur near an orbital phase of 0.5, and about 95%/99% of them reside within the phase range of $\sim 0.43\text{--}0.58/\sim 0.3\text{--}0.72$, respectively. The two most extreme secondary phases in our catalog are for TIC 149673382, with a secondary phase of ≈ 0.87 , and TIC 337097515, with a secondary phase of ≈ 0.11 , both exhibiting a pronounced heartbeat “bump” between the primary and secondary eclipses.

We note that several targets in our catalog have primary depths larger than 0.5 according to the current version of ELEANOR data, likely due to systematics. The three most extreme cases are TIC 42066695 (average depth of ≈ 848 ppt), TIC 446208053 (depth of ≈ 804 ppt), and TIC 192305147 (depth of ≈ 751 ppt); the phase-folded lightcurves for the first two are shown in Figure 21. The average primary depth for 42066695 is much smaller, ≈ 434 ppt; for the other two targets there is no publicly available QLP data at the time of writing.

It is important to note that the observed eclipse depths often vary from one sector to the next. With a handful of exceptions, mentioned below, these depth variations are due to systematic effects inherent to the lightcurve extraction process. An example is shown in Figure 22 for the case of TIC 5232381, where the eclipses in Sector 9 (left panel) are about half as deep as those in Sector 62 (right panel). This is likely due to the sector-specific background subtraction being affected by TIC 5232374, $\approx 13''$ away and about 2 mag fainter.

TIC 5232381 is neither an isolated occurrence nor an outlier. Sometimes, eclipse depths can fluctuate even within a

single sector, showing differences before and after TESS data downlink gaps. In the most extreme cases, the eclipses can be virtually undetectable in certain sectors, as highlighted in Figure 23 for TIC 77392704 (also, e.g., for TIC 63165670). Thus, in these cases it is preferable to exclude such sectors when phase-folding the lightcurve, which we did by visual inspection on a target-by-target and sector-by-sector basis.

These complications are often further exacerbated when the number of detected eclipses is small due to relatively long orbital periods, data gaps, and systematic effects. Sometimes, even the addition of new sectors does not help improve the measurements. An example of this is shown in Figure 24 for the case of TIC 173706211, where the Sector 84 lightcurve is completely dominated by systematics and there is a single useful eclipse near the end of the sector. As a result, while obtaining reliable eclipse depths and durations for individual sectors is, in general, relatively straightforward, extending these measurements across multiple sectors is challenging and sector-averaged depths and durations can be misleading.

Collectively, these factors underscore the complexity of obtaining accurate and precise measurements for many TESS FFI EBs. Thus, it is important to emphasize that even after thorough scrutiny it is still possible that the ephemerides provided in this catalog are slightly off, especially when the number of TESS observations is small, the eclipses are few and shallow, the SNR is low, and the lightcurve is dominated by systematics. Unfortunately, resolving these issues by, e.g., cross-checking EBs between different lightcurve pipelines is far from straightforward. For instance, it is not uncommon to see depth differences when comparing ELEANOR data to QLP or TESS-SPOC data, likely due to different treatments of crowding correction.⁴⁰

Finally, it is worth pointing out two further considerations regarding eclipse depth variations in particular. First, these do

⁴⁰ We note that, at the time of writing, the QLP lightcurves beyond Sector 74 were not well suited for ephemeris determination due to a timing error of about 3 minutes. One of us (D. Short) noticed this issue during our investigations and brought it to the attention of the QLP team.




Figure 25. The distributions of Gaia’s AEN, AENS, RUWE, and effective temperatures for the 7936 new EBs presented here. The vertical dashed lines represent AENS = 3 (left panel) and RUWE = 1.4 (middle panel), potentially suggesting unresolved companions.

not substantially affect the ephemeris measurements as these are done on an eclipse-by-eclipse basis—each eclipse is independently modeled with a generalized Gaussian that fits for the depth by design. Neither do the depth variations dramatically impact our photocenter-based vetting as it mostly depends on the eclipse durations.⁴¹

5.2. EBs in Multiple Stellar Systems

Multiple stellar systems are not uncommon. About one in 10 binary stars reside in hierarchical (2+1) triples, and thousands of even higher order systems have already been discovered (e.g., D. Raghavan et al. 2010; S. Tremaine 2020; and references therein). The higher the multiplicity of the system, the higher its complexity in terms of orbital and physical parameters, formation and evolution pathways, and long-term dynamical stability (e.g., M. Moe & R. Di Stefano 2017; S. Tremaine 2020; A. Tokovinin 2021; and references therein). In general, wide multiple systems are prime targets for long-term astrometric monitoring, while compact multiples are ideally suited for observing short-term dynamical interactions between their components such as ETVs.

Crossmatching our catalog with the Gaia DR3 astrometric measurements, we extracted the available `astrometric_excess_noise` (AEN), `astrometric_excess_noise_sig` (AENS), and renormalized unit weight error (RUWE). These can be used to test for unseen companions (e.g., V. Belokurov et al. 2020; Z. Penoyre et al. 2020; K. G. Stassun & G. Torres 2021; P. Gandhi et al. 2022; S. R. Majewski et al. 2025, in preparation; and references therein), which, if indeed present, would potentially mark the EBs as components in systems of three (or more) stars. The corresponding distributions are shown in Figure 25, highlighting several interesting features. In particular, the AEN is greater than 10 mas for hundreds of targets, and the AENS is greater than 3 for ~40% of the EBs, reaching values of tens to even hundreds of thousands for dozens of targets. Similarly, the RUWE is greater than 1.4—suggesting unresolved companions (K. G. Stassun & G. Torres 2021)—for about one in every four targets. Altogether, these considerations indicate that a potentially large fraction of the 7936 new EBs presented here may reside in multiple stellar systems.

Another option for finding multiple stellar systems is through the presence of extra events in the lightcurves of

EBs. Indeed, TESS has already enabled the detection of thousands of such events, practically revolutionizing the field by discovering hundreds of new 2+1 triply eclipsing triple systems (e.g., T. Borkovits 2022; S. A. Rappaport et al. 2022, 2024; V. B. Kostov et al. 2024b; and references therein) and 2+2 eclipsing quadruple systems (V. B. Kostov et al. 2022b, 2024a; P. Zasche et al. 2024; B. P. Powell et al. 2025), as well as unusual (2+1)+1 eclipsing quadruples (e.g., B. P. Powell et al. 2022b), several (2+1)+2 quintuple systems (V. B. Kostov et al. 2022b, 2024a), the first two (2+2)+2 eclipsing sextuple systems (B. P. Powell et al. 2021a; P. Zasche et al. 2023), and even two transiting circumbinary planets (V. B. Kostov et al. 2020, 2021a). Volunteers at EBP have independently rediscovered many of these and, naturally, a significant number of false positives that mimic 2+2 eclipsing quadruples due to blended light from two unrelated EBs,⁴² and also identified several new eclipsing triple and quadruple candidates (V. B. Kostov 2025, in preparation).

5.3. Known EBs Observed by TESS

Given our ML search was effectively blind, it was inevitable that it picked up a large number of known EBs. And indeed, as discussed above, about one in four of the identified candidates are within 1 pixel of known EBs. Thus, in order to verify the efficiency and reliability of our automated ephemeris and vetting pipeline, we applied it to ~30,000 such targets and tracked its performance. Interestingly, during the early stages of the EBP project, the volunteers noticed that the correctly measured periods from TESS were sometimes different from the literature values. Altogether, we marked 2065 such cases. As an example, the distributions of the period ratios between TESS on the one hand and Gaia, ASAS-SN, ATLAS, and VSX on the other are highlighted in Figure 26.

As seen from the figure, most of the Gaia, ASAS-SN, ATLAS, and VSX periods are close to an integer fraction of the true period, where for simplicity “close” is defined as within 10% of integer fractions of 2 (from 1/2 to 10/2) and 3 (from 1/3 to 20/3).⁴³ This is perhaps not too surprising given the much longer continuous baseline coverage and higher cadence of TESS observations compared to other surveys. As an example, Figure 27 shows TIC 2597145, where the correct period measured from TESS is 1.4143 days, twice the period

⁴¹ Naturally, it would be important to keep track of such effects when validating targets based on eclipse depths and contamination from nearby stars.

⁴² See, for example, <https://www.zooniverse.org/projects/vbkostov/eclipsing-binary-patrol/talk/tags/multiple-system-candidate>.

⁴³ The reciprocal fractions are combined in the figure, i.e., 1/2 with 2/1, 1/3 with 3/1, 2/3 with 3/2, etc.




Figure 26. Distributions of ratios between the correct TESS periods and the incorrect periods from Gaia (first panel from left), ASAS-SN (second panel from left), ATLAS (third panel from left), and VSX (last panel from left) for 2065 known EBs. For simplicity, the reciprocal fractions are combined, i.e., $1/2$ with $2/1$, $1/3$ with $3/1$, etc. The zero values represent targets of which the Gaia/ASAS-SN/ATLAS/VSX periods are not within 10% of a corresponding integer fraction of the TESS periods.




Figure 27. Left panels: ephemeris and vetting pipeline results for TIC 2597145. The correct period measured from TESS is 1.4143 days—twice as long as the period listed in Gaia (0.7072 days); ASAS-SN and VSX provide the correct period. Right panels: same as left but for TIC 9473243. Here, the Gaia period is 4 times the correct period, the ASAS-SN period is twice the correct period, and the WISE period is half of the correct period.




Figure 28. Same as Figure 27 but for TIC 403072759 (left) and TIC 12396863 (right). Thanks to TESS, shallow secondary eclipses can be detected near phase 0.5 for both targets, confirming the corresponding periods are 1.3029 days and 2.6760 days. For comparison, the periods listed in Gaia are 0.5210 days for TIC 403072759 (i.e., two-fifths of the true period) and 5.3519 days for TIC 12396863 (i.e., twice the true period).

listed in Gaia (0.7072 days). For this target, the periods listed in ASAS-SN and VSX are correct. Another example is TIC 9473243, where the correct period measured from TESS is 2.2669 days, whereas Gaia gives a period of 9.0680 days (4 times as long), ASAS-SN gives a period of 4.5338 days (2 times as long), and WISE gives a period of 1.1335 days (only half as long). Additionally, TESS excels at enabling the detection of shallow secondary eclipses. An example of this is shown in Figure 28 for TIC 403072759, highlighting the shallow but clear secondary eclipse near the phase of 0.5. Here, the true period measured from TESS is 1.3029 days

whereas Gaia gives a period of 0.52 days, i.e., a $2/5$ fraction of the true period.

Interestingly, about 20% of the Gaia EB periods seem to be unrelated to the TESS periods at all. These cases are represented in Figure 26 by the peak at zero. One example is TIC 2239760, where the correct TESS period is 5.8855 days, while the Gaia period is 3.2370 days, a ratio of ≈ 0.55 (Figure 29, left panel). Another is TIC 143060048, where the TESS period is 4.2852 days and the Gaia period is 30.4715 days (ratio of ≈ 7.11 ; Figure 29, right panel). Some of the most extreme discrepancies are for TIC 443450339, 152328270,




Figure 29. Same as Figure 27 but for TIC 2239760 (left) and TIC 143060048 (right), for which the corresponding ratios between the Gaia and TESS periods are not close to low-order integer ratios (≈ 0.55 and ≈ 7.11 , respectively).




Figure 30. Upper panel: TESS FFI ELEANOR lightcurve of TIC 307119043, an eclipsing 2+2 quadruple exhibiting two sets of primary and secondary eclipses (V. B. Kostov et al. 2022b). Lower panel: eclipsing (2+1)+1 quadruple system TIC 114936199 exhibiting a complex eclipse on the outer orbit (B. P. Powell et al. 2022b).

353628656, 138032974, and 34853800, for which the TESS periods are 2.9032, 9.0706, 2.5816, 1.8713, and 3.2028 days, respectively, whereas the corresponding Gaia periods are orders of magnitude longer, i.e., 406.0884, 381.0741, 355.6257, 189.4914, and 164.4351 days.

Table 4 highlights 10 random rows of the catalog of 2065 known EBs with updated ephemerides produced as part of this work.

5.4. Interesting Systems

Here, we highlight some of the more interesting targets independently identified as part of this effort, split into the following categories:

1. Additional eclipses: Targets exhibiting extra events not associated with the EB signal, such as a second set of

eclipses following a different period (representing a 2+2 quadruple system consisting of two EBs) or complex tertiary events (representing triply eclipsing 2+1 triple systems or eclipsing (2+1)+1 quadruple systems). Figure 30 highlights two such examples.

2. ETVs: Targets where the eclipse times deviate from the linear ephemeris, suggesting potential dynamical interactions with additional bodies. An example is shown in the upper panel of Figure 31 for the case of the 2+2 quadruple system TIC 219006972, where the two EB subsystems are dynamically interactive on observable timescales (V. B. Kostov et al. 2023). Another example is the known EB TIC 26542657 (A. Prša et al. 2022), which, through our comparison with the Kepler catalog, we determined does not show the ~ 11 day eclipses in Kepler and must therefore be a higher-order system. It is




Figure 31. Upper panel: primary ETVs for the two EBs in the 2+2 quadruple system TIC 219006972, confirming the two subsystems are gravitationally bound (V. B. Kostov et al. 2023) with an outer period of 168 days. Lower panel: primary (red) and secondary (blue) ETVs of the known EB TIC 26542657, suggesting a 2 +1 triple system with an outer period of about 300 days.




Figure 32. TESS ELEANOR lightcurve of the known EB TIC 26542657, highlighting clear changes in the shape of the primary and secondary eclipses between the first two sectors (14 and 15; sharp primary, rounded secondary) and the last three sectors (81, 82, and 83; rounded primary, flat secondary), along with a prominent tertiary eclipse in Sector 81 (red arrow).




Figure 33. Kepler lightcurve of the known TESS EB TIC 26542657 showing a single tertiary eclipse (marked with a red arrow) and no discernible eclipses from the ~ 11 day EB.




Figure 34. Disappearing eclipses due to dynamical interactions with unseen companions. Upper panel: TIC 236774836 (T. Mitnyan et al. 2024b); lower panel: TIC 220410224.




Figure 35. TESS lightcurve of the known EB TIC 234229841, where the primary and secondary eclipses “switch” places. In Sector 6, the deeper eclipses precede the heartbeat-like hump, the primary and secondary eclipses have similar depths in Sector 33, and in Sector 87 the deeper eclipses follow the hump.




Figure 36. Pronounced apsidal motion exhibited by TIC 189281140 (upper panel) and TIC 470715046 (lower panel).

separated by only $\sim 1''$ from TIC 1882992210, which is only ~ 0.1 mag fainter in TESS, making photocenter confirmation of the eclipse source effectively impossible from TESS. However, the target exhibits clear primary and secondary ETVs (see Figure 31, lower panel), a tertiary eclipse in Sector 81, and prominent changes in the shape of both the primary and secondary eclipses

between Sectors 14/15 (narrow, sharper primary, more rounded secondary) and later sectors (more rounded primary, flat secondary; see Figure 32). Taken together, these provide strong evidence that either TIC 26542657 or TIC 1882992210 is a dynamically interacting, triply eclipsing triple system with an outer period of about 300 days. Additionally, as seen from Figure 33, the Kepler




Figure 37. Two sectors of ELEANOR data for TIC 21159577, an EB showcasing an evolving pattern of starspot-induced rotational modulations.




Figure 38. Upper panel: TESS lightcurve of an EB with a pulsating component (TIC 22621932). Lower panel: TESS lightcurve of an EB exhibiting a pronounced heartbeat pattern (TIC 336538437; S. Solanki et al. 2025).

lightcurve of the target shows one tertiary eclipse suggesting that the system was out of the eclipsing window for the ~ 11 day EB during the Kepler era due to orbital precession.⁴⁴

3. Disappearing eclipses: Targets where the detected eclipses exhibit prominent depth variations due to precession of the EB orbital plane, to the point of eventually ceasing altogether. Figure 34 shows the TESS lightcurves of TIC 236774836 (T. Mitnyan et al. 2024b) and TIC 220410224, indicating dynamical interaction with unseen companions.
4. “Switching” eclipses: Similar to the previous example, but here the depth ratio between the primary and

secondary eclipses changes between sectors. Figure 35 shows an example of this effect for the case of the known EB TIC 234229841.

5. Apsidal motion: Targets exhibiting pronounced “smear” of the secondary eclipses in orbital phase, indicating apsidal motion. Figure 36 highlights two such targets, TIC 189281140 and TIC 470715046.
6. Stellar variability: Targets exhibiting prominent light-curve modulations due to, e.g., rotational variability (spotted stars), pulsating components, heartbeat patterns, etc. Figures 37 and 38 show examples of each category, represented by TIC 21159577, TIC 22621932, and TIC 336538437.
7. Transiting planets: It is only logical that a search for stellar eclipses will result in finding planetary transits as well. Indeed, our ML pipeline picked up the confirmed planet TIC 408310006 (WASP-166 b). Interestingly, an

⁴⁴ We note that the tertiary eclipse has nonzero quality flags, which were not uncommon in Kepler data, and could be easily missed if one only investigates “quality = 0” data.




Figure 39. Four sectors of short-cadence TESS data for WASP-166 (TIC 408310006) showing the prominent transits of the known planet WASP-166 b. Two additional transit-like events can be seen in Sectors 62 and 89 (highlighted with red arrows), suggesting the potential presence of a second planet in the system.

eagle-eyed volunteer on EBP (D. Iannone) noticed an additional transit-like event in the lightcurve of the target in Sector 62.⁴⁵ Further investigation showed another event in Sector 89 (see Figure 39), suggesting the potential presence of a second transiting planet in the system.

6. Summary

We have presented the TESS Ten Thousand Catalog containing 10,001 uniformly vetted and validated EBs observed by TESS in FFI data. Of these 7936 are new EBs while the remaining 2065 are known EBs of which the period listed in one or more catalogs is incorrect. The targets were detected by a neural network search applied to the Sector 1 through 26 lightcurves. These were produced with a local implementation of the ELEANOR pipeline, and extracted for all stars brighter than TESS magnitude $T = 15$. The EBs passed comprehensive automated analysis and thorough visual scrutiny by citizen scientists, including confirmation of the measured ephemerides and photocenter offsets, and cross-matching against millions of known EBs from multiple catalogs. Most of the 7936 new EBs are on the fainter end (median magnitude $T = 13.8$), have short orbital periods (median period of 3.5 days), have eclipses that originate within

$\sim 1''$ – $2''$ of the respective TIC star, and have been observed in at least three TESS sectors. For the 2065 known EBs, we corrected the ephemerides available at the time of writing. Astrometric measurements from Gaia suggest that a significant fraction of the new EBs may have unresolved companions and thus be part of higher-order stellar systems. In addition, some of the new EBs show ETVs, apsidal motion, and even extra eclipses due to additional stars. These are excellent targets for further in-depth investigation aimed at unraveling the underlying architecture and dynamics. Finally, we provide a list of $\sim 900,000$ unvetted and unvalidated TESS targets for which the neural network identified eclipse-like features and scored them higher than 0.9, and for which there are no known EBs within a sky-projected separation of 1 TESS pixel ($21''$).

Acknowledgments

This paper includes data collected by the TESS mission, which are publicly available from the Mikulski Archive for Space Telescopes (MAST). Funding for the TESS mission is provided by NASA’s Science Mission Directorate.

This research has made use of the Exoplanet Follow-up Observation Program website, which is operated by the California Institute of Technology, under contract with the National Aeronautics and Space Administration under the Exoplanet Exploration Program.

This work has made use of data from the European Space Agency (ESA) mission Gaia (<https://www.cosmos.esa.int/web/gaia>), processed by the Gaia Data Processing and

⁴⁵ <https://www.zooniverse.org/projects/vbkostov/eclipsing-binary-patrol/talk/6324/3432679?comment=5638772&page=1>

Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the Gaia Multilateral Agreement.

This publication uses data generated via the Zooniverse.org platform, the development of which is funded by generous support, including a Global Impact Award from Google, and by a grant from the Alfred P. Sloan Foundation.

The EBP project was made possible by the participation of nearly 2000 incredible citizen scientists who volunteered their free time for the advancement of our understanding of EBs. The project as a whole is referred to as the Eclipsing Binary Patrol Collaboration in the affiliations listed above.

Resources supporting this work were provided by the NASA High-end Computing (HEC) Program through the NASA Center for Climate Simulation (NCCS) at Goddard Space Flight Center.

Resources supporting this work were provided by the NASA HEC Program through the NASA Advanced Supercomputing Division at Ames Research Center for the production of the SPOC data products.

V.B.K. is grateful for financial support from NASA grant 80NSSC21K0631; V.B.K., J.O., and W.W. acknowledge support from NSF grant AST-2206814.

Facilities: Gaia, MAST, TESS.

Software: Exogram (<https://exogram.vercel.app/>), Fast Lightcurve Inspector (<https://fast-lightcurve-inspector.osc-fr1.scalingo.io/>), astropy (Astropy Collaboration et al. 2013, 2018, 2022), Keras (F. Chollet 2015), Lightcurve (Lightcurve Collaboration et al. 2018), Matplotlib (J. D. Hunter 2007), NumPy (C. R. Harris et al. 2020), Pandas (W. McKinney 2010), SciPy (P. Virtanen et al. 2020), TensorFlow (M. Abadi et al. 2015).

ORCID iDs

Veselin B. Kostov  <https://orcid.org/0000-0001-9786-1031>
 Brian P. Powell  <https://orcid.org/0000-0003-0501-2636>
 Aline U. Fornear  <https://orcid.org/0000-0001-7062-7632>
 Marco Z. Di Fraia  <https://orcid.org/0000-0001-7701-6818>
 Robert Gagliano  <https://orcid.org/0000-0002-5665-1879>
 Thomas L. Jacobs  <https://orcid.org/0000-0003-3988-3245>
 Julien S. de Lambilly  <https://orcid.org/0000-0003-3464-1554>
 Hugo A. Durantini Luca  <https://orcid.org/0000-0002-4143-2550>
 Steven R. Majewski  <https://orcid.org/0000-0003-2025-3147>
 Jerome Orosz  <https://orcid.org/0000-0001-9647-2886>
 Saul A. Rappaport  <https://orcid.org/0000-0003-3182-5569>
 Donald Short  <https://orcid.org/0000-0001-5504-9512>
 William Welsh  <https://orcid.org/0000-0003-2381-5301>
 Svetoslav Alexandrov  <https://orcid.org/0000-0002-4959-8598>
 Cledison Marcos da Silva  <https://orcid.org/0000-0001-8953-9149>
 Erika Dunning  <https://orcid.org/0009-0001-2901-8875>
 Michiharu Hyogo  <https://orcid.org/0000-0001-8343-0820>
 Christian Magliano  <https://orcid.org/0000-0001-6343-4744>
 John Yablonsky  <https://orcid.org/0000-0002-0248-4817>
 Sovan Acharya  <https://orcid.org/0000-0002-7759-0569>
 Fred Adams  <https://orcid.org/0000-0002-8167-1767>
 Thomas Barclay  <https://orcid.org/0000-0001-7139-2724>

Benjamin T. Montet  <https://orcid.org/0000-0001-7516-8308>

Susan Mullally  <https://orcid.org/0000-0001-7106-4683>

Greg Olmschenk  <https://orcid.org/0000-0001-8472-2219>

Andrej Prša  <https://orcid.org/0000-0002-1913-0281>

Elisa Quintana  <https://orcid.org/0000-0003-1309-2904>

Robert Wilson  <https://orcid.org/0000-0002-4235-6369>

Ethan Kruse  <https://orcid.org/0000-0002-0493-1342>

References

- Abadi, M., Agarwal, A., Barham, P., et al., 2015 TensorFlow: Large-scale machine learning on heterogeneous systems, v2.19, <https://www.tensorflow.org/>
- Abdul-Masih, M., Prša, A., Conroy, K., et al. 2016, *AJ*, **151**, 101
- Andersen, J. 1991, *A&ARv*, **3**, 91
- Armstrong, D. J., Gamper, J., & Damoulas, T. 2021, *MNRAS*, **504**, 5327
- 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
- Bahdanau, D., Cho, K., & Bengio, Y. 2014, arXiv:1409.0473
- Belokurov, V., Penoyre, Z., Oh, S., et al. 2020, *MNRAS*, **496**, 1922
- Borkovits, T. 2022, *Galax*, **10**, 9
- Borkovits, T., Mitnyan, T., Rappaport, S. A., et al. 2022, *MNRAS*, **510**, 1352
- Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Sci*, **327**, 977
- Boyajian, T. S., LaCourse, D. M., Rappaport, S. A., et al. 2016, *MNRAS*, **457**, 3988
- Burke, C. J., Levine, A., Fausnaugh, M., et al., 2020 TESS-Point: High precision TESS pointing tool, v0.9, Astrophysics Source Code Library, ascl: 2003.001
- Cacciapuoti, L., Kostov, V. B., Kuchner, M., et al. 2022, *MNRAS*, **513**, 102
- Caldwell, D. A., Tenenbaum, P., Twicken, J. D., et al. 2020, *RNAAS*, **4**, 201
- Capistrant, B. K., Soares-Furtado, M., Vanderburg, A., et al. 2022, *ApJS*, **263**, 14
- Chollet, F., 2015 Keras, v3.10.0, <https://keras.io>
- Conroy, K. E., Prša, A., Stassun, K. G., et al. 2014a, *AJ*, **147**, 45
- Conroy, K. E., Prša, A., Stassun, K. G., et al. 2014b, *PASP*, **126**, 914
- Coughlin, J. L., Thompson, S. E., Bryson, S. T., et al. 2014, *AJ*, **147**, 119
- Eggen, O. J. 1957, *Obs*, **77**, 191
- Eisner, N., 2022 LATTE: Lightcurve Analysis Tool for Transiting Exoplanet, Astrophysics Source Code Library, ascl:2205.006
- Eisner, N. L., Barragán, O., Lintott, C., et al. 2021, *MNRAS*, **501**, 4669
- Fang, X., Thompson, T. A., & Hirata, C. M. 2018, *MNRAS*, **476**, 4234
- Feinstein, A. D., Montet, B. T., Foreman-Mackey, D., et al. 2019, *PASP*, **131**, 094502
- Fischer, D. A., Schwamb, M. E., Schawinski, K., et al. 2012, *MNRAS*, **419**, 2900
- Fragione, G., & Kocsis, B. 2019, *MNRAS*, **486**, 4781
- Gandhi, P., Buckley, D. A. H., Charles, P. A., et al. 2022, *MNRAS*, **510**, 3885
- Gao, X., Chen, X., Wang, S., & Liu, J. 2025, *ApJS*, **276**, 57
- Goodfellow, I., Bengio, Y., & Courville, A. 2016, *Deep Learning* (Cambridge, MA: MIT Press)
- Goodricke, J. 1973, *RSPT*, **73**, 474
- Green, M. J., Maoz, D., Mazeh, T., et al. 2023, *MNRAS*, **522**, 29
- Hamers, A. S., Rantala, A., Neunteufel, P., Preece, H., & Vynatheya, P. 2021, *MNRAS*, **502**, 4479
- Han, T., & Brandt, T. D. 2023, *AJ*, **165**, 71
- Handler, G., Kurtz, D. W., Rappaport, S. A., et al. 2020, *NatAs*, **4**, 684
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
- Hartman, J. D., Bakos, G. Á., Bouma, L. G., & Csubry, Z. 2025, *PASP*, **137**, 024501
- He, K., Zhang, X., Ren, S., & Sun, J. 2016, in 2016 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (Piscataway, NJ: IEEE), 770
- Heinze, A. N., Tonry, J. L., Denneau, L., et al. 2018, *AJ*, **156**, 241
- Hochreiter, S., & Schmidhuber, J. 1997, *Neural Comput.*, **9**, 1735
- Howard, E. L., Davenport, J. R. A., & Covey, K. R. 2022, *RNAAS*, **6**, 96
- Huang, C. X., Vanderburg, A., Pál, A., et al. 2020a, *RNAAS*, **4**, 204
- Huang, C. X., Vanderburg, A., Pál, A., et al. 2020b, *RNAAS*, **4**, 206
- Hunter, J. D. 2007, *CSE*, **9**, 90
- IJspeert, L. W., Tkachenko, A., Johnston, C., & Aerts, C. 2024, *A&A*, **691**, A242

- Jayaraman, R., Rappaport, S. A., Powell, B., et al. 2024, *ApJ*, **975**, 121
- Kirk, B., Conroy, K., Prša, A., et al. 2016, *AJ*, **151**, 68
- Kochanek, C. S. 2021, *MNRAS*, **507**, 5832
- Kopal, Z. 1956, *AnAp*, **19**, 298
- Kostov, V. B. 2023, *Univ*, **9**, 455
- Kostov, V. B., Borkovits, T., Rappaport, S. A., et al. 2023, *MNRAS*, **522**, 90
- Kostov, V. B., Kuchner, M. J., Caciapiuti, L., et al. 2022a, *PASP*, **134**, 044401
- Kostov, V. B., Mullally, S. E., Quintana, E. V., et al. 2019, *AJ*, **157**, 124
- Kostov, V. B., Orosz, J. A., Feinstein, A. D., et al. 2020, *AJ*, **159**, 253
- Kostov, V. B., Powell, B. P., Orosz, J. A., et al. 2021a, *AJ*, **162**, 234
- Kostov, V. B., Powell, B. P., Rappaport, S. A., et al. 2022b, *ApJS*, **259**, 66
- Kostov, V. B., Powell, B. P., Rappaport, S. A., et al. 2024a, *MNRAS*, **527**, 3995
- Kostov, V. B., Powell, B. P., Torres, G., et al. 2021b, *ApJ*, **917**, 93
- Kostov, V. B., Rappaport, S. A., Borkovits, T., et al. 2024b, *ApJ*, **974**, 25
- Kotikalapudi, R., 2017 keras-vis, v0.8.7, GitHub, <https://github.com/raghakot/keras-vis>
- Kovács, G., Zucker, S., & Mazeh, T. 2002, *A&A*, **391**, 369
- Kozai, Y. 1962, *AJ*, **67**, 591
- Kristiansen, M. H. K., Rappaport, S. A., Vanderburg, A. M., et al. 2022, *PASP*, **134**, 074401
- Kunimoto, M., Bryson, S., Daylan, T., et al. 2024, AAS Meeting, **243**, 231.05
- Kunimoto, M., Tey, E., Fong, W., et al. 2022, *RNAAS*, **6**, 236
- LaCourse, D. M., Jek, K. J., Jacobs, T. L., et al. 2015, *MNRAS*, **452**, 3561
- Lea, C., Flynn, M. D., Vidal, R., Reiter, A., & Hager, G. D. 2017, in 2017 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (Piscataway, NJ: IEEE), 156
- LeCun, Y., Boser, B., Denker, J. S., et al. 1989, *Neural Comput.*, **1**, 541
- Lidov, M. L. 1962, *P&SS*, **9**, 719
- Lightcurve Collaboration, Cardoso, J. V. d. M., Hedges, C., et al., 2018 Lightcurve: Kepler and TESS time series analysis in Python, v2.5, Astrophysics Source Code Library, ascl: **1812.013**
- Liu, B., & Lai, D. 2019, *MNRAS*, **483**, 4060
- Luong, M.-T., Pham, H., & Manning, C. D. 2015, Proc. of the 2015 Conf. on Empirical Methods in Natural Language Processing (Association for Computational Linguistics), 1412
- Magliano, C., Kostov, V., Caciapiuti, L., et al. 2023, *MNRAS*, **521**, 3749
- Matijević, G., Prša, A., Orosz, J. A., et al. 2012, *AJ*, **143**, 123
- McKinney, W. 2010, Proc. 9th Python in Science Conf., ed. S. van der Walt & J. Millman, (SciPy), 51
- Melton, E. J., Feigelson, E. D., Montalto, M., et al. 2024, *AJ*, **167**, 203
- Mitnyan, T., Borkovits, T., Czavalinga, D. R., et al. 2024a, *A&A*, **686**, C3
- Mitnyan, T., Borkovits, T., Czavalinga, D. R., et al. 2024b, *A&A*, **685**, A43
- Moe, M., & Di Stefano, R. 2017, *ApJS*, **230**, 15
- Montalto, M. 2023, *MNRAS*, **518**, L31
- Mowlavi, N., Holl, B., Lecoœur-Taïbi, I., et al. 2023, *A&A*, **674**, A16
- Nair, V., & Hinton, G. E. 2010, in Proc. of the 27th Int. Conf. on Machine Learning (ICML-10) (Madison, WI: Omnipress), 807
- Nelson, S. 2008, *Natur*, **455**, 36
- Niemela, V. 2001, *RMxAC*, **11**, 23
- Oelkers, R. J., & Stassun, K. G. 2018, *AJ*, **156**, 132
- Offner, S. S. R., Moe, M., Kratter, K. M., et al. 2023, in ASP Conf. Ser. 534, Protostars and Planets VII, ed. S. Inutsuka et al. (San Francisco, CA: ASP), 275
- Oláh, K., Seli, B., Haris, A., et al. 2025, *A&A*, **698**, A150
- Osterbrock, D. E. 1953, *ApJ*, **118**, 529
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *JMLR*, **12**, 2825
- Pejcha, O., Antognini, J. M., Shappee, B. J., & Thompson, T. A. 2013, *MNRAS*, **435**, 943
- Penoyre, Z., Belokurov, V., Wyn Evans, N., Everall, A., & Koposov, S. E. 2020, *MNRAS*, **495**, 321
- Petrosky, E., Hwang, H.-C., Zakamska, N. L., Chandra, V., & Hill, M. J. 2021, *MNRAS*, **503**, 3975
- Powell, B. P., Kostov, V. B., Rappaport, S. A., et al. 2021a, *AJ*, **161**, 162
- Powell, B. P., Kostov, V. B., Rappaport, S. A., et al. 2021b, *AJ*, **162**, 299
- Powell, B. P., Kostov, V. B., & Tokovinin, A. 2023, *MNRAS*, **524**, 4296
- Powell, B. P., Kruse, E., Montet, B. T., et al. 2022a, *RNAAS*, **6**, 111
- Powell, B. P., Rappaport, S. A., Borkovits, T., et al. 2022b, *ApJ*, **938**, 133
- Powell, B. P., Torres, G., Kostov, V. B., et al. 2025, *ApJ*, **985**, 213
- Prša, A., Batalha, N., Slawson, R. W., et al. 2011, *AJ*, **141**, 83
- Prša, A., Kochoska, A., Conroy, K. E., et al. 2022, *ApJS*, **258**, 16
- Raghavan, D., McAlister, H. A., Henry, T. J., et al. 2010, *ApJS*, **190**, 1
- Rappaport, S., Vanderburg, A., Jacobs, T., et al. 2018, *MNRAS*, **474**, 1453
- Rappaport, S. A., Borkovits, T., Gagliano, R., et al. 2022, *MNRAS*, **513**, 4341
- Rappaport, S. A., Borkovits, T., Gagliano, R., et al. 2023, *MNRAS*, **521**, 558
- Rappaport, S. A., Borkovits, T., Mitnyan, T., et al. 2024, *A&A*, **686**, A27
- Ricker, G. R., Winn, J. N., Vanderspek, R., et al. 2015, *JATIS*, **1**, 014003
- Rowan, D. M., Jayasinghe, T., Stanek, K. Z., et al. 2022, *MNRAS*, **517**, 2190
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. 1986, *Natur*, **323**, 533
- Russell, H. N. 1948, *HarMo*, **7**, 181
- Shan, Y., Chen, J., Zhang, Z., et al. 2025, *PASP*, **137**, 044503
- Shara, M. M., Howell, S. B., Furlan, E., et al. 2021, *MNRAS*, **507**, 560
- Shi, X., Chen, Z., Wang, H., et al. 2015, Advances in Neural Information Processing Systems 28, ed. C. Cortes et al. (NeurIPS) <https://proceedings.neurips.cc/paper/2015/hash/07563a3fe3bbe7e3ba84431ad9d055af-Abstract.html>
- Simonyan, K., & Zisserman, A. 2014, arXiv:1409.1556
- Slawson, R. W., Prša, A., Welsh, W. F., et al. 2011, *AJ*, **142**, 160
- Solanki, S., Cieplak, A. M., Schnittman, J., et al. 2025, *ApJS*, **276**, 17
- Soszyński, I., Udalski, A., Szymański, M. K., et al. 2017, *AcA*, **67**, 297
- Stassun, K. G., Oelkers, R. J., Paegert, M., et al. 2019, *AJ*, **158**, 138
- Stassun, K. G., & Torres, G. 2021, *ApJL*, **907**, L33
- Sullivan, P. W., Winn, J. N., Berta-Thompson, Z. K., et al. 2015, *ApJ*, **809**, 77
- Sun, C., Shrivastava, A., Singh, S., & Gupta, A. 2017, in 2017 IEEE Int. Conf. on Computer Vision (ICCV) (Piscataway, NJ: IEEE), 843
- Thiemann, H. B., Norton, A. J., Dickinson, H. J., McMaster, A., & Kolb, U. C. 2021, *MNRAS*, **502**, 1299
- Thompson, S. E., Coughlin, J. L., Hoffman, K., et al. 2018, *ApJS*, **235**, 38
- Thompson, S. E., Mullally, F., Coughlin, J., et al. 2015, *ApJ*, **812**, 46
- Tokovinin, A. 2021, *Univ*, **7**, 352
- Torres, G., Andersen, J., & Giménez, A. 2010, *A&ARv*, **18**, 67
- Trani, A. A., Rastello, S., Di Carlo, U. N., et al. 2022, *MNRAS*, **511**, 1362
- Tremaine, S. 2020, *MNRAS*, **493**, 5583
- Twicken, J. D., Catanzarite, J. H., Clarke, B. D., et al. 2018, *PASP*, **130**, 064502
- Valizadegan, H., Martinho, M. J. S., Jenkins, J. M., et al. 2025, arXiv:2502.09790
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
- Vynatheya, P., & Hamers, A. S. 2022, *ApJ*, **926**, 195
- Wang, J., Fischer, D. A., Barclay, T., et al. 2015, *ApJ*, **815**, 127
- Welsh, W. F., & Orosz, J. A. 2018, in Handbook of Exoplanets, ed. H. J. Deeg & J. A. Belmonte (Cham: Springer), 34
- Zasche, P., Borkovits, T., Jayaraman, R., et al. 2023, *MNRAS*, **520**, 3127
- Zasche, P., Henzl, Z., Merc, J., Kára, J., & Kučáková, H. 2024, *A&A*, **687**, A6
- Zeipel, H. V. 1910, *AN*, **183**, 345