**Publ. Astron. Obs. Belgrade No. 100 (2021), 275 - 280** *Progress Report*

# COMPUTER VISION AS A TOOL FOR STUDYING CLOSE BINARY STARS

OLIVERA LATKOVIĆ and ATILA ČEKI
*Astronomical Observatory, Volgina 7, 11000 Belgrade, Serbia*
*E–mail: olivia@aob.rs*

**Abstract.** Computer vision is a subfield of artificial intelligence that deals with automated detection and classification of objects from images, used in a variety of advanced applications, from facial recognition to self-driving cars. While other machine learning methods have gained a strong footing in astronomy over the last several years, computer vision is still a relatively rare and novel approach. We have been experimenting with this technique in the context of eclipsing binary stars, with the aim to automate the analysis of photometric time-series data from ground- and space-based surveys. The current and future deluge of such data requires the automation of as many tasks as possible, otherwise much of it will remain unutilized. Computer vision might be used to estimate the stellar and orbital parameters of eclipsing binaries based on the images of their light curves. As proof-of-concept, we have already developed a computer vision system for automated recognition of light curves with total eclipses, and demonstrated that a computer can perform this task as well or better than humans.

## 1. INTRODUCTION

Close binaries, although well studied, are still at the forefront of astrophysics because they allow us to determine the masses, sizes and temperatures of stars based on relatively simple physics. In eclipsing systems, these parameters are traditionally found by light curve modeling. This involves fitting synthetic light curves, generated from a mathematical model parametrized directly by the orbital parameters of the system and the stellar parameters of the components, to the observations. Several such models are in use. The most popular is the so-called “WD Code” (see e.g. Wilson et al. 2020 and the references within), but a quantity of excellent work has also been done with the model developed within our institute, the “Djurasevic code” (Djurašević 1992, Djurašević et al. 1998) which has recently been modernized for automated applications (see e.g. Djurašević et al. 2016, Latković et al. 2019). This is the model we use for the research reported here.

Finding accurate orbital and stellar parameters of an eclipsing binary from a model is possible because they are encoded in the shape of the light curve. For example, the widths and relative depths of eclipses are indicative of the temperatures and sizes of the components, the amplitude of the light curve is related to the orbital inclination, and so forth. However, a critical parameter, the mass ratio, typically cannot be measured from photometry alone, and requires the construction of radial velocity curves, which are in turn obtained from time-series spectroscopy. In the case of contact binaries, where the components are immersed in a stable, co-rotating common envelope, spectroscopy can be skipped. The shape of the common envelope is determined by a single equipotential Roche surface, which means the mass ratio is strictly constrained by the size ratio, and the size ratio can be obtained from light curve modeling—but only in the presence of a total eclipse (see the contribution of Čeki et al. in these proceedings for a more substantial discussion).

With the work presented here, we aim to automate the selection of totally eclipsing contact binaries from vast light curve archives of ongoing and future ground- and space-based surveys using a machine learning technique—computer vision. Automation of this and other tasks related to the modeling of close binaries is getting more and more important as the amounts of underutilized observations grow. For example, some 30 000 contact binaries were identified in the Catalina Real-Time Transient Survey (CRTS) Variable Star Catalog (Drake et al. 2014); and more than 70 000 in the ASAS-SN Catalog of Variable Stars (Jayasinghe et al. 2020). These samples are already too large to tackle "manually", by traditional examination of individual objects; and yet even larger ones are sure to follow from Gaia, LSST and so forth.

A human astronomer recognizes a total eclipse as a distinct, flat region inside the light minimum by looking at a plot of the light curve. Computer vision encompasses machine learning algorithms that essentially do the same: recognize objects in images. In what follows, we describe the architecture, training and performance of a computer vision system for detection of total eclipses in the light curves of contact binary stars.

## 2. COMPUTER VISION MODELS FOR AUTOMATED RECOGNITION OF TOTAL ECLIPSES

Computer vision methods are usually based on convolutional neural networks, or CNNs. A CNN reduces the information contained in a matrix of image pixels to a numeric array that can be processed by an ordinary classification algorithm through a series of convolutions with detection filters—low-dimension matrices trained to select visual features such as vertical or horizontal edges, curves and so on.

Like other supervised machine learning algorithms, CNNs learn by processing large numbers of labeled examples. In our case, that would be the images of light curves with either partial or total eclipses. *Unlike* most other machine learning algorithms, CNNs don't need precalculated features to help them discriminate between objects in images. If we wanted to apply an ordinary neural network to light curves as time series, we would have to engineer the features ourselves. This was done in one of the first applications of machine learning in astronomy, for automatic modeling of thousands of eclipsing binaries observed by Kepler: the light curves were first fitted by chain polynomials, and the coefficients of these polynomials were used as features for the neural network (Prša et al. 2011). A CNN learns the features itself by finding statistical patterns in the data.

The simplest of our CNN models has three convolutional layers that transform input images of 32×32 pixels into 2×2×32 volumes, and eventually feature vectors with 128 elements that are fed to a classification layer that resembles an ordinary neural network. Increasingly complex models, with 4, 5 or 6 convolutions, are trained to process larger images (up to 512×512 pixels), with full light curves or eclipse cutouts.

### 2. 1. GENERATED TRAINING DATA

For training, we use synthetic light curves created using the Djurasevic code. Some examples are given in Fig. 1. We generate equal numbers of light curves with total and partial eclipses. Labels for the generated images (whether the eclipse is total or not) are assigned automatically based on the light contribution of the eclipsed star. The model parameters for generating these light curves (orbital inclination, mass ratio, the size of the primary star relative to orbit size, its effective temperature and the parameters of a single spot) are drawn randomly from distributions derived using a large sample of individually studied W UMa stars (see the contribution of Čeki et al. in these proceedings).

We generate training images with three levels of noise: noiseless (similar to binned Kepler light curves; see Fig. 1), moderate (representative of what we expect from ground-based observations of individual objects) and noisy (similar to what we have seen in the light curves from ASAS, OGLE and CRTS). The noise is drawn from a normal distribution with the standard deviation approximately equal to 2 and 5% of the light curve amplitude in the latter two cases, respectively.

### 2. 2. TEST DATA AND HUMAN PERFORMANCE

We also use synthetic light curves to validate the model performance by testing it on labeled data that were not used for training, but we test the model on real observations as well. However, here we face another problem: that of the reliability of human classifications and in general, the availability of ground truth.

With generated data we have access to the fluxes of both stars and can determine the true label beyond doubt. But even simulated total eclipses can be made to look like partial ones in a noisy light curve or if there are too few points in the minimum to make out its shape.

To our knowledge, human ability to recognize total eclipses has never been questioned or measured prior to this work. We measured it on synthetic data (for which we could later look up the ground truth). The two of us (Č.A. and L.O.) looked at randomized sets of 500 generated images in two sizes[1] with all three noise levels, and got the results summarized in Table 1. On average, we achieved 97% balanced accuracy[2] on noiseless examples, down to 85% on noisy examples. Meaning that, in the presence of noise, we misclassified 15% of light curves.

Let us see how this compares to computer vision.

[1]This experiment was done in an early stage of the research and the images used for it are of different sizes (250×250 and 350×175 pixels; see also Table 1) than those we currently use (32×32, 64×64, 128×128 pixels, 256×256 and 512×512 pixels). The 250×250 pixel images are not significantly different from the 256×256 pixel or other square images. The 350×175 pixel images are twice as wide as they are tall, which is helpful (for humans) when trying to determine the nature of the eclipses, but has proven to have no notable effect on the performance of the machine learning method. This is why we discontinued their use. The measurement of human performance with updated image shapes and sizes (and hopefully with more human participants) is part of planned future work.

[2]Balanced accuracy is a performance metric often used instead of the ordinary accuracy (that would be the ratio of correct classifications to the sample size) to account for the "class imbalance" expected in observational data. Namely, there are far more light curves with partial than with total eclipses. In the case of Kepler catalog, out of about 300 contact binary light curves, only around 30 (or 10%) show total eclipses (at least as far as we humans can tell). While not as severe as in many other applications of machine learning, this imbalance can affect the performance of our eclipse recognition system in ways we're yet to quantify.

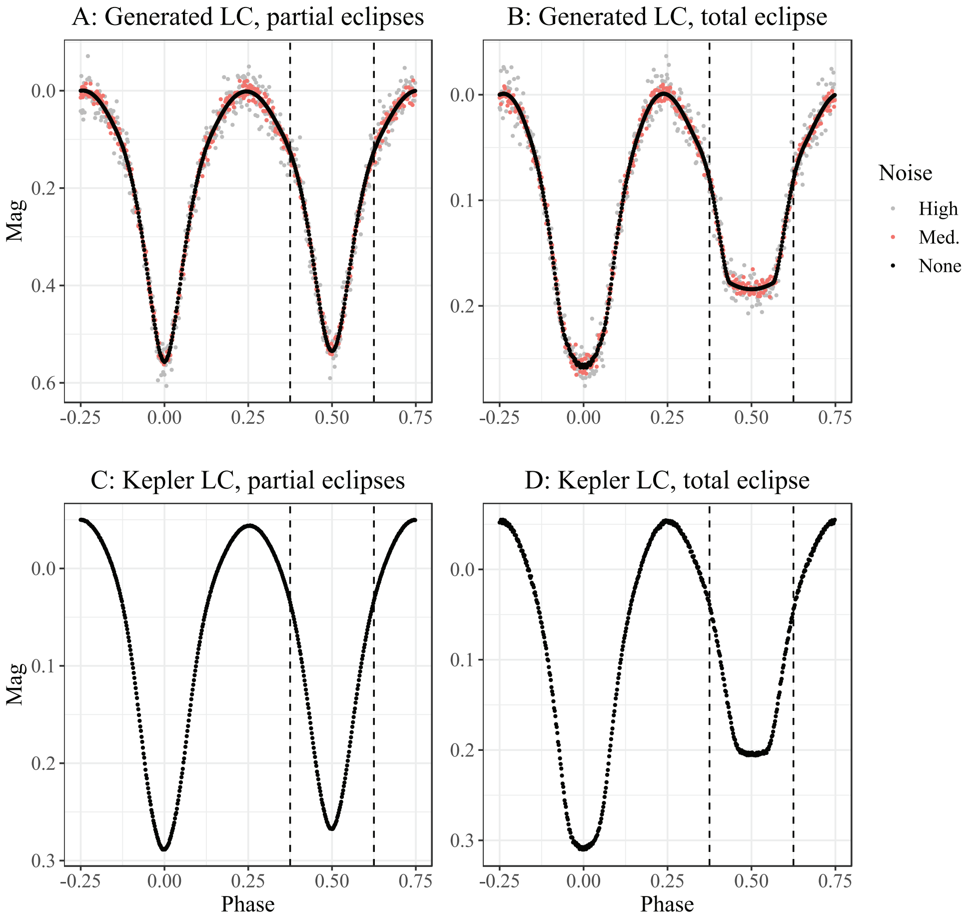


Figure 1: Top – examples of generated light curves with different noise levels and partial (panel A) or total eclipses (panel B). Bottom – examples of binned Kepler light curves, again with partial (panel C, showing data for KIC 10074939) or total eclipses (panel D, showing data for KIC 5022908). These examples are representative of the largest input images (512×512 pixels) used for training and testing of the computer vision algorithm. The vertical dashed lines indicate the phase range for input images with eclipse cutouts (128×128 pixels and smaller).

### 2. 3. PERFORMANCE OF COMPUTER VISION

Fig. 2A shows how our models fared when evaluated with test data generated the same way as training data. The colors represent different noise levels and the performance obviously declines as the noise increases. The different line types represent images of different sizes; increasing image size moderately improves performance. Finally, from left to right is the size of the training set, starting at 500 and ending at 75 000 training examples. Each set is divided equally between the two classes. The dotted black lines mark the human performance levels from Table 1 for each noise level. Apparently, AI can surpass humans with this task—at least with synthetic light curves.

Table 1: Human performance

| Noise level | Image size | Balanced accuracy |
|---|---|---|
| | 250x250 | 0.96 |
| | 350x175 | 0.97 |
| **0.00** | | **0.97** |
| | 250x250 | 0.88 |
| | 350x175 | 0.92 |
| **0.02** | | **0.90** |
| | 250x250 | 0.84 |
| | 350x175 | 0.85 |
| **0.05** | | **0.85** |

In Fig. 2B we see the same plot, but when the models are evaluated using binned Kepler observations instead. The best models, trained with moderate noise, reach just above 90% balanced accuracy and surpass human performance on noisy light curves. The models trained on noisy data perform poorly, but this is expected, because the training light curves generated with high noise are visually very different from smoothed Kepler observations.

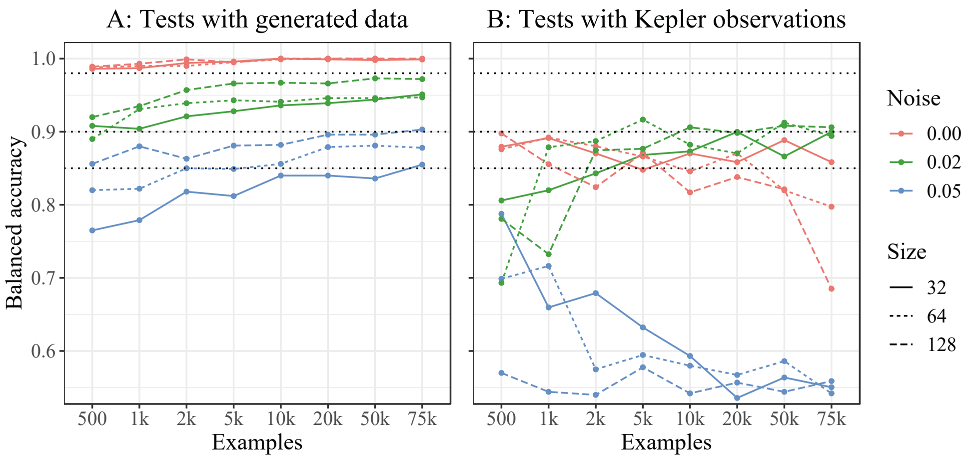


Figure 2: Performance of computer vision in recognition of total and partial eclipses when tested with generated data (panel A) and with binned observations from Kepler (panel B). The horizontal dotted lines mark the human performance at different noise levels (see Table 1).

## 3. CONCLUSIONS AND FUTURE WORK

The success of training a machine learning algorithm using simulated data depends critically on the ability of the simulation to mimic real-life phenomena the algorithm is intended to work with *after* training. In our case, great performance on generated light curves but significantly worse on real observations means that there's a "data mismatch". To improve this situation, we might need to re-examine our assumptions about the nature of the noise in the observations, and take into account the factual class imbalance when creating simulated training sets.

There are also other datasets we can test our models on. The application on ASAS data is already a work in progress. Preliminary tests indicate that models trained on noisy generated data perform the best there, but overall, not as well as with Kepler. Another venue is CRTS data, where Sun et al. (2020) have already identified 3000 totally eclipsing systems among the 30 000 light curves of contact binaries.

Once we've established the reliability of our computer vision system, it can be applied to even larger, existing and future datasets, as a first step of a fully automated analysis pipeline.

## Acknowledgments

The research presented in this report was funded by the Ministry of Education, Science and Technological Development of the Republic of Serbia (contract No. 451-03-68/2020-14/200002).

## References

Djurašević, G.: 1992, *Astrophys. & Sp. Sc.*, **196**, 241.
Djurašević, G., Zakirov, M., Hojaev, A., Arzumanyants, G.: 1998, *Astron. Astrophys. Suppl.*, **131**, 17.
Djurašević, G., Essam, A., Latković, O., Čeki, A., El-Sadek, M. A., Abo-Elala, M. S., Hayman, Z. M.: 2016, *Astron. J.*, **152**, A57.
Drake, A. J., Graham, M. J., Djorgovski, S. G. et al.: 2014, *Astrophys. J. Suppl. Series*, **213**, 9.
Jayasinghe, T., Stanek, K. Z., Kochanek, C. S. et al.: 2020, *Mon. Not. R. Astron. Soc.*, **493**, 4045.
Latković, O., Čeki, A., Djurašević, G., Essam, A,, Hamed, A. S., Youssef, S. M.: 2019, *Astron. J.*, **157**, 3.
Prša, A, Batalha, N, Slawson, R. W. et al.: 2011, *Astron J.*, **141**, 83.
Sun, W., Chen, X., Deng, L., de Grijs, R.: 2020, *Astrophys. J. Suppl. Series*, **247**, 50.
Wilson, R. E., Devinney, E. J., and Van Hamme, W.: 2020, *Astrophysics Source Code Library.*