

Detection of exomoons in simulated light curves with a regularized convolutional neural network

Rasha Alshehhi¹, Kai Rodenbeck^{2,3}, Laurent Gizon^{3,2,1}, and Katepalli R. Sreenivasan^{4,1}

¹ Center for Space Science, NYUAD Institute, New York University Abu Dhabi, Abu Dhabi, UAE
e-mail: ra130@nyu.edu

² Institut für Astrophysik, Georg-August-Universität Göttingen, Friedrich-Hund-Platz 1, 37077 Göttingen, Germany

³ Max-Planck-Institut für Sonnensystemforschung, Justus-von-Liebig-Weg 3, 37077 Göttingen, Germany

⁴ Tandon School of Engineering, New York University, New York, USA

Received 5 November 2019 / Accepted 12 May 2020

ABSTRACT

Context. Many moons have been detected around planets in our Solar System, but none has been detected unambiguously around any of the confirmed extrasolar planets.

Aims. We test the feasibility of a supervised convolutional neural network to classify photometric transit light curves of planet-host stars and identify exomoon transits, while avoiding false positives caused by stellar variability or instrumental noise.

Methods. Convolutional neural networks are known to have contributed to improving the accuracy of classification tasks. The network optimization is typically performed without studying the effect of noise on the training process. Here we design and optimize a 1D convolutional neural network to classify photometric transit light curves. We regularize the network by the total variation loss in order to remove unwanted variations in the data features.

Results. Using numerical experiments, we demonstrate the benefits of our network, which produces results comparable to or better than the standard network solutions. Most importantly, our network clearly outperforms a classical method used in exoplanet science to identify moon-like signals. Thus the proposed network is a promising approach for analyzing real transit light curves in the future.

Key words. Moon – methods: numerical – methods: data analysis

1. Introduction

In the past two decades, over 3500 planets around distant stars (exoplanets) have been detected and confirmed. Most of the known exoplanets have been detected using the transit method, in which space telescopes observe a star for a period of time, generating a photometric light curve (Rodenbeck et al. 2018). If an exoplanet passes in front of the star on its orbit around the star, light is blocked out and the brightness decreases. Planetary transits repeat every orbital period. The shape of a transit light curve, in combination with additional knowledge about the host star, is used to constrain the characteristics of the planet. If an exoplanet has a moon, it is also expected to leave a signature in the transit light curve. Detecting exomoons is difficult because a stellar variability component as well as photon noise are present (Rodenbeck et al. 2018).

Several methods have been proposed to detect exomoons around exoplanets in light curves, using individual transits or averages over multiple transits (Kipping 2009; Heller 2014). A planet-moon transit light curve is modeled simply as the sum of two transits; one for the planet, and a more shallow transit for the moon. To model multiple planet-moon transits, more detailed modeling is required in principle, to take the orbital motion of the moon and of the planet around their centers of mass into account. This more detailed modeling is only applied when a promising target has been identified. For previous attempts for detecting exomoons, we refer to Szabó et al. (2013) and Hippke (2015). The most recent attempt has been reported by Teachey et al. (2018), who stacked the phase-folded

transits of 284 *Kepler* exoplanets in search for exomoon candidates.

Unfortunately, no exomoon has been unambiguously detected around any exoplanet so far, although some exomoon candidates have been reported. The most promising such candidate orbits the exoplanet *Kepler* 1625 b. This candidate has a mass of 10–40 Earth masses (M_{\oplus}) and a radius of $4.90^{+0.79}_{-0.72}$ Earth radii (R_{\oplus}), which is higher and larger than the masses and radii of many confirmed exoplanets. In total, four transits of *Kepler*-1625 b have been observed: three by the *Kepler* space telescope between 2009 and 2012, and one from the *Hubble* space telescope in October 2017 (Teachey et al. 2018; Teachey & Kipping 2018). *Kepler*-1625 b orbits its star every 287 days; the exomoon orbits the planet with a period between 13 and 39 days (see Fig. 1 for a diagram of the orbital configuration).

We propose a solution based on a convolutional neural network (CNN) to detect exomoons in photometric light curves. To test the method, we generate synthetic transit light curves with and without an exomoon. The light curves are similar to those of the *Kepler*-1625 b observations. Figure 2 presents examples of such simulated light curves. The left panel shows a planetary transit including an exomoon at low noise level. The right panel shows a transit containing an exomoon at higher noise level. The exomoon transit is always close in space to the planet transit because the exomoon is on a close-in orbit of the planet. Depending on the orbital configuration, the exomoon transit can occur before or after the planetary transit.

Recently, deep CNN architectures have achieved impressive performance for data classification. They are quickly becoming

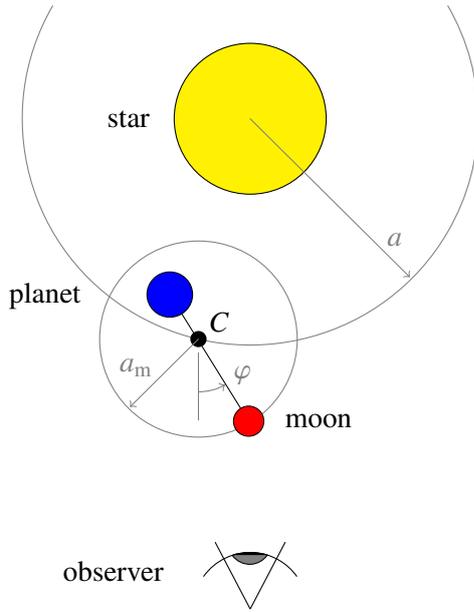


Fig. 1. Sketch of the orbital configuration for the numerical simulations of transit light curves. The planet and the moon orbit their common center of mass C .

prominent in astrophysical applications because they have the ability to effectively encode spectral and spatial information based on the input data, without pre-processing. A typical network consists of multiple interconnected layers and learns a hierarchical feature representation from raw data (Liu & Deng 2015). The network optimization is usually performed with L2-norm without considering the effect of the noise properties, and it does not exploit neighboring information in the data (Janocha & Czarnecki 2017). We here propose a 1D CNN to classify a transit signal as either containing or excluding an exomoon candidate. The architecture consists of five 1D convolution layers, followed by one fully connected layer. Gaussian smoothness is used as total variation loss to penalize the L1-norm error. The goal of using this smoothness is to remove the noise or unwanted variations in the data (e.g., from stellar activity, or systematic or instrumental effects) and conserve the neighborhood information at the same time.

Deep learning has led to significant advances in the field of astrophysics. Most of the studies that incorporate this method were based on well-known network configurations, such as alexnet (Krizhevsky et al. 2012) and vgg16 (Liu & Deng 2015), which is tuned using a mean standard error as cost function. We here propose an improved cost function that includes a total variation loss component to regularize the network. Total variation loss has been used extensively in other fields, for example, in computer vision. Mahendran & Vedaldi (2015) used the total variation as a regularization procedure for reconstructing natural images from their representation based on image priors. Javanmardi et al. (2016) used the total variation as a structured loss by applying a Sobel edge-detection technique on the output probability map to train a deep network, especially in the case when there are not enough labeled data. We apply similar ideas to improve the classification of exoplanet light curves.

The remainder of this paper is organized as follows. In Sect. 2 we describe the data sets we used in the analysis. An overview of our proposed method is given in Sect. 3. Section 4 compares the performance of the proposed method with the performances of

several other techniques in both generalization and convergence. Section 5 summarizes the most important findings.

2. Synthetic data

We used three sets of light curves (Table 1). Each light curve contained four transits. For each transit we had 50 days of data, with 25 days on each side of a transit. The light curves had a cadence of 29.4 min, which is *Kepler's* long cadence. The simulated transit light curves were generated using the model described in detail by Rodenbeck et al. (2018), with the additional improvement that the occultation of the moon by the planet was taken into account.

The first data set consists of 1 052 400 simulated light curves; 526 200 with exomoons and 526 200 without exomoons. The noise was varied between 25 ppm and 500 ppm. The radius was also varied between $0.25 R_{\oplus}$ to $5 R_{\oplus}$. The exomoon phase was randomized between 0 and 1 for each light curve that contained a moon.

The second data set consisted of ten subsets. Each subset had 200 000 curves: 100 000 with exomoons and 100 000 without exomoons. The noise was kept constant at 100 ppm. The exomoon radius was varied from $0.25 R_{\oplus}$ to $4.75 R_{\oplus}$ from one subset to another. The exomoon phases were also randomized.

The third data set also contained ten subsets, each with 200 000 curves. There was either no moon (100 000) or a moon (100 000) with a radius of $1 R_{\oplus}$. The noise was varied from 25 ppm to 475 ppm. The exomoon phases were again randomized.

Table 2 summarizes the choice of astrophysical parameters for the simulations, that is, the stellar, planetary, and exomoon parameters. The stellar parameters were fixed; the star had the same radius R_{\star} and mass M_{\star} as our Sun. We used a quadratic limb-darkening law, which describes the dimming of a star's brightness from the center of the stellar disk toward the limb (see Claret & Bloemen 2011 for more detail), with roughly solar-like values. The exoplanet radius R_p and mass M_p are roughly Jupiter-like, and the orbital distance was 0.85 astronomical units. The exoplanet orbits its star in 287 days, in accordance with Kepler's third law. This law describes the relation between the distance of a planet to its star, the orbital period of the planet, and the combined mass of the star, the planet, and the moon if present, to determine the distance from the star. The exomoon semimajor axis around its planet a_m was set to $20 R_{\text{Jupiter}}$. The planet mass and the exomoon semimajor axis determine the orbital period of the moon around its planet. We varied the exomoon radius between 0.25 and $5 R_{\oplus}$ while keeping its mass density constant and equal to the Earth's density. We chose to keep the density constant, therefore only one exomoon parameter was varied. We also added noise of various amplitudes between 25 ppm (parts per million) and 500 ppm to the synthetic light curves, instead of using a more realistic noise model. The initial exomoon phase $\varphi_0 = \varphi(t = 0)$ (see Fig. 1) was randomized for each realization. The realization means a new generation of the light curve with the appropriate parameters, with newly generated noise added to the model. The q_1 and q_2 determine the brightness of the star as a function of center-to-limb distance. Limb darkening affects the shape of the transit: the bottom of the transit is rounded and not flat (see Fig. 2).

3. Method

In this section, the proposed framework for detection of a moon-like signal from a light curve is described. The method does not

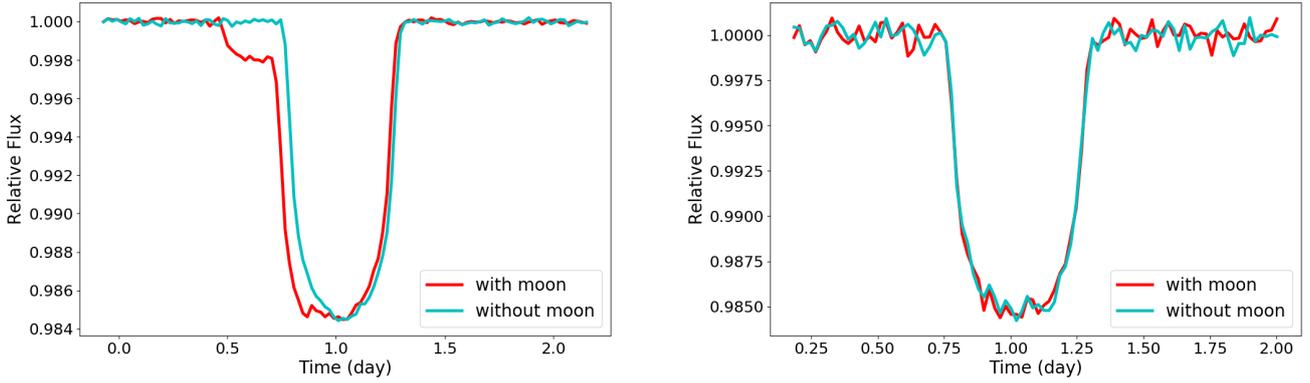


Fig. 2. Examples of simulated transit light curves. *Left panel:* the curve shows a transit of a planet with radius $R_p = 12 R_\oplus$ (similar to that of Kepler 1625b) for a low level of noise (100 ppm). The red curve shows a similar but different planetary transit, and an exomoon transit is superimposed ($R_m = 4.75 R_\oplus$). *Right panel:* simulations for a higher noise level (400 ppm) and a much smaller moon $R_m = R_\oplus$. The distinction between the planet-only (blue) and planet-moon (red) case is not visible by eye in most cases.

Table 1. Parameters for the three data sets of simulated light curves.

Data set	Exomoon radius	Noise level	M
1	$0.25\text{--}5 R_\oplus$	25–500 ppm	1 052 400
2	$0.25\text{--}4.75 R_\oplus$	100 ppm	$10 \times 200\,000$
3	$1 R_\oplus$	25–475 ppm	$10 \times 200\,000$

Notes. The second and third columns give the range of values for the exomoon radius and the noise level, respectively. The last column gives the total number of simulated light curves, M .

Table 2. Stellar, planetary, and exomoon parameters for the transit light-curve simulations.

Parameter	Value (or range)
Planet-star radius ratio, r_p	0.11
Barycenter semimajor axis, a_b	$183 R_\star$
Barycenter impact parameter b	0
Barycenter initial orbital phase φ_b	0
Barycenter orbital period P_b	287 days
Stellar limb-darkening parameters	$q_1 = 0.6, q_2 = 0.4$
Moon-star radius ratio, r_m	0.00229–0.04584
Moon semimajor axis, a_m	$2.055 R_\star$
Moon initial orbital phase, φ_m	$0 - 2\pi$ (random)
Moon orbital period, P_m	0.5523 days
Moon-planet mass ratio, f_m	1.2×10^{-5} –0.0942

Notes. The subscripts \odot and \oplus refer to the Sun and Earth values, respectively.

require a preprocessing step. The 1D CNN architecture is summarized below and described in full in Appendix A. Our definition of the loss function is presented below.

The idea of detecting a moon in light curves is similar to the idea of detecting an edge in images, where brightness (pixel value) changes sharply. The convolution stage in CNNs involves the sharing of connection weights, and this allows for robust position-invariant feature detection. Time variations are detected regardless of when they happen to be within the window of interest.

The input of the CNN is a simulated light curve, and the output is a binary value indicating the presence of a moon. We denote the m th simulated light curve by $X_m = \{X_m(0), X_m(1) \dots X_m(N-1)\}$,

where $0 \leq m \leq M-1$. Each light curve contains four transits. The total number of time samples is $N = 9796$, and the number of simulated light curves, M , is given in Table 1. For each input simulated light curve, the output of the network is denoted by $C(X_m)$, equal to 0 (no moon) or 1 (moon). The true classification of the light curve is $c_m = 1$ (moon) or $c_m = 0$ (no moon). We note that convolutional networks are also appropriate because they are computationally efficient; the dimension of each input is reduced by convolution and pooling stages.

3.1. Convolutional neural network architecture

Our network takes X_m as input and classifies it into a binary value. As presented in Fig. 3, the architecture of the network consists of five convolution layers, followed by one fully connected layer and one output layer. Each convolution layer consists of a 1D convolution, batch normalization, leaky ReLU activation, maximum pooling, and Gaussian dropout. The determination of the network architecture was obtained experimentally.

3.2. Loss function

The training process of the CNN is accomplished through the minimization of a loss function that measures the error between correct and predicted values. In the classification problems, there are many loss functions that could be used (e.g., binary cross-entropy, hinge, Kullback–Leibler (KL) divergence, and Wasserstein) (Janocha & Czarnecki 2017).

Noise-smoothing operators have been extensively used as preprocessing in the context of signal processing to suppress the noise and preserve the changes at the same time (Simonoff 1998). The contribution of our work is the use of the smoothing operator directly in learning. The loss function is constructed using spatial information (e.g., neighborhood information) of the data features. As a result, back-propagation errors of the proposed loss functions are different. In this section, we define the proposed loss function that combines mean error and total variation error.

We consider a binary classification, where the goal is the assignment of either no-moon (0) or moon (1) to each moon-like signal in the training set. The input data are a set of training samples indexed by $0 \leq m \leq M-1$. The loss function is a combination of two functions. The first function is the mean error,

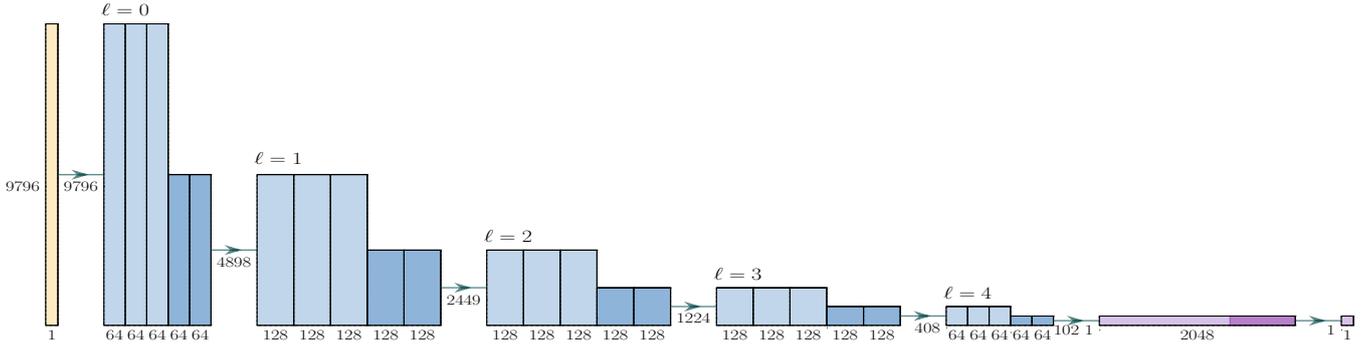


Fig. 3. CNN architecture. As explained in the text, it consists of five convolution layers composed of 64, 128, 128, 128, and 64 channels, including convolution, batch normalization, leaky ReLU activation (first three columns in light blue), maximum pooling, and Gaussian dropout (last two columns in dark blue). The last layer is followed by a fully connected layer (purple) and a classification layer returning a binary answer (1 for moon, 0 otherwise).

$$L_{\text{err}}(\Theta) = \frac{1}{M} \sum_{m=0}^{M-1} |C(X_m; \Theta) - c_m|, \quad (1)$$

where each term in the sum can only take the values 1 or 0. This function is to be minimized to obtain all the parameters Θ of the CNN (all weights and biases). However, if the number of parameters Θ is large and the amount of data is low, the problem might be overfit. Regularization by penalizing the loss function is a common solution for this problem.

The output of the last convolution layer ($\ell = 4$) is a matrix of dimension 64 (number of feature maps) times 102 (number of time samples). This layer is followed by a fully connected layer $Y(p)$ activated by a leaky ReLU (see Appendix). We use the total variation of Y as a penalty to regularize the classification. First, each vector Y_m is smoothed to remove noise while preserving the original signal. This is done by convolution with a centered Gaussian kernel G with standard deviation $\sigma = 3$ and width 9,

$$S_m(p) = \sum_{j=0}^8 Y_m(p-j)G(j). \quad (2)$$

The total variation loss is then obtained by computing the mean unsigned difference between two consecutive points, and averaging over all the light curves,

$$L_{\text{tv}}(\Theta) = \frac{1}{M} \sum_{m=0}^{M-1} \frac{1}{N_p - 1} \sum_{p=0}^{N_p-2} |(S_m(p+1) - S_m(p))|. \quad (3)$$

The final loss function is a linear combination of the mean error between the correct and predicted values and the average unsigned gradient of the input features,

$$L(\Theta) = L_{\text{err}}(\Theta) + \gamma L_{\text{tv}}(\Theta), \quad (4)$$

where γ is a regularization parameter. We here give equal weight to the two terms, that is, $\gamma = 1$.

3.3. Optimization

In order to obtain the network parameters Θ^* that minimize $L(\Theta)$, we use the AdamMax algorithm to gradually update the weights and the biases when searching for the optimal solution, see Appendix B.

4. Performance

We present evaluation metrics and setup (Sect. 4.1) to assess the performance of the CNN in comparison to the classical method for the detection of exomoons (Sect. 4.2). In Sect. 4.3 the performance of the proposed method is further discussed by varying the parameters of the moon, the planet, and the host star.

4.1. Performance metrics and setup

The most common metrics for the evaluation of a binary classification are sensitivity, specificity, and accuracy. Sensitivity, also called the true-positive rate, measures the proportion of actual positive samples that are correctly classified as positive samples. Specificity, also called true-negative rate, is the proportion of actual negative samples that are correctly classified as negative examples (Hand et al. 2001). Accuracy measures the proportion of correctly classified light curves. These metrics are defined as follows:

$$\text{Sen} = \frac{TP}{TP + FN}, \quad (\text{sensitivity}) \quad (5)$$

$$\text{Spe} = \frac{TN}{TN + FP}, \quad (\text{specificity}) \quad (6)$$

$$\text{Acc} = \frac{TP + TN}{TP + FP + FN + TN}, \quad (\text{accuracy}) \quad (7)$$

where TP is the number of true positives (exomoons correctly detected), FP is the number of false positives, TN is the number of true negatives (absence of moons correctly classified), and FN is the number of false negatives.

We used 70% of each data set (in Table 1) to train and validate the CNN and 30% of data to evaluate the CNN performance using the above metrics. All our CNN experiments were run on Nvidia Tesla V100 GPUs with Keras 2.2.2 with TensorFlow 1.9.1. Each CNN experiment with the same configuration was run 25 times. The results we present here were obtained by averaging over these 25 trials. The code of the network can be found online¹.

4.2. Classical analysis

For comparison with the CNN algorithm, we also classified the light curves using the classical method described in Rodenbeck et al. (2018). For each light curve, we considered

¹ <https://github.com/ralshehhi/Exomoon>

Table 3. Number of epochs needed to train the model and computing time in second required to process the test data for different choices of loss functions.

Loss function	Training (number of epochs)	Testing (Computation time (s))
Binary entropy	24	149.7
Hinge	100	148.9
KL	100	148.3
Wasserstein	100	145.7
L_{err}	20	146.8
$(L_{\text{err}} + L_{\text{tv}})$	23	146.8

Notes. Each epoch consists of 32 samples from data set 1.

two models that might explain the data: a light-curve model containing the transits of a single planet without a moon (which we call no-moon model), and a model containing the transits of both a planet and a moon (which we call one-moon model). To estimate the likelihood of each case, we used the Bayesian information criterion (BIC, Schwarz 1978) to determine whether a moon is present. The BIC of a model $\mathcal{M}(\theta)$, which depends on model parameters $\theta = \{\theta_1, \theta_2 \dots \theta_k\}$, is given by

$$\text{BIC}(\mathcal{M}) = -2 \max_{\theta} \mathcal{L}(\theta|X) + k \ln N, \quad (8)$$

where $\mathcal{L}(\theta|X)$ is the likelihood of θ given the data X , k is the number of used parameters ($k = 14$ for the one-moon model and $k = 7$ for the no-moon model), and $N = 9796$ is the length of X , as described in Sect. 3. We calculated the difference $\Delta\text{BIC} = \text{BIC}(\text{one moon}) - \text{BIC}(\text{no moon})$ between the two BICs and classified light curves with a $\Delta\text{BIC} > 0$ as containing no moon and $\Delta\text{BIC} < 0$ as containing a moon. The ΔBIC between two competing models compared the maximum likelihood of the two models while penalizing the model with more parameters. In other words, the model with more parameters has to explain the data sufficiently better to justify the use of more free parameters. The one-moon model contains five more parameters than the no-moon model, which describes the moon size and orbital configuration (e.g., ratio of the moon-to-star radius, planet-moon distance, moon period, moon phase, and ratio of moon-to-planet mass). We find the maximum likelihood of the model with a moon and without a moon using the Markov chain Monte Carlo (MCMC) sampler *emcee*, which is used to approximate the parameter posterior distribution $\mathcal{L}(\theta|X)$.

4.3. Results and analysis

4.3.1. Effect of the choice of loss function

In order to illustrate the effectiveness of the total variation loss function L defined in Sect. 3.2, we considered other loss functions: binary cross-entropy, hinge, KL, wasserstein, and mean square error functions. Tables 3 and 4 show faster convergence and better performance of L than all other choices of loss functions.

4.3.2. Effect of smoothing as preprocessing

The main purpose of the preprocessing step is to produce more effective features by standardizing the dynamic range of the raw data or to remove unwanted variations in the raw data. Smoothing is one of the preprocessing steps that might be used to remove unwanted variations. However, this approach could

Table 4. Comparison between loss functions based on performance metrics: sensitivity, specificity, and accuracy.

Loss function	Sen (%)	Spe (%)	Acc (%)
Binary entropy	97	49	73
Hinge	100	0	50
KL	100	0	50
Wasserstein	0	100	50
L_{err}	90	96	93
$(L_{\text{err}} + L_{\text{tv}})$	94	97	96

Notes. Data set 1 was used.

Table 5. Comparison between smoothing techniques.

Method	Sen (%)	Spe (%)	Acc (%)
Median + L_{err}	81	99	90
Cosine + L_{err}	92	90	91
Hamming + L_{err}	93	81	87
Gaussian + L_{err}	92	88	90
No prefilter + $L_{\text{err}} + L_{\text{tv}}$	94	97	96

Notes. The first four rows show performance results obtained by applying smoothing (with different kernels) before the CNN using L_{err} to train the network. The last row presents the performance metrics of CNN with the loss function (Eq. (4)) without smoothing. Data set 1 was used.

introduce new systemic variability or remove actual signal from the raw data. In Table 5 we compare the smoothing method as a preprocessing step of 1D CNN with L_{err} and the method we propose here, which uses L_{tv} to learn the parameters.

4.3.3. Dependence on exomoon radius

Figure 4 shows the results we obtained by increasing the exomoon radius. Sen and Spe are significantly lower when the exomoon radius is smaller than the Earth radius, although the convolution, pooling, and smoothing windows are small enough to capture small differences in data distribution. When the exomoon radius is equal to or greater than the Earth radius, the performance reaches 100%. The total variation is commonly used as a regularization factor to find continuity in the signals (Houhou et al. 2009) and detect a change between neighboring values, but the variations due to small moons are very difficult to identify.

The addition of a small moon alone alters the light curve slightly, which also means that the maximum likelihoods of the no-moon and one-moon models are very similar. In this case, the ΔBIC method decides in favor of the model with the fewer parameters, and Sen goes to zero.

4.4. Effect of noise on the light curves

Figure 5 shows the results for detecting a $1 R_{\oplus}$ moon when the amount of noise is varied from 25 ppm to 475 ppm every 50 ppm. We find that when the noise is lower than 175 ppm, the performance of the network is perfect, and it is very good (above 80%) when the noise is below 375 ppm. The performance decreases

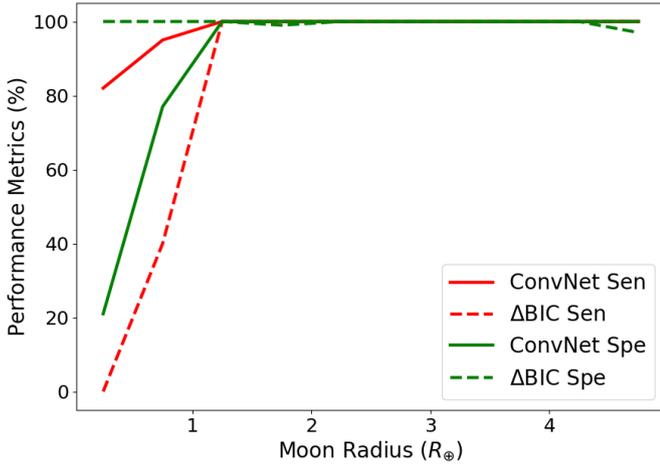


Fig. 4. Performance as a function of exomoon radius at fixed noise level (100 ppm). Data set 2 was used.

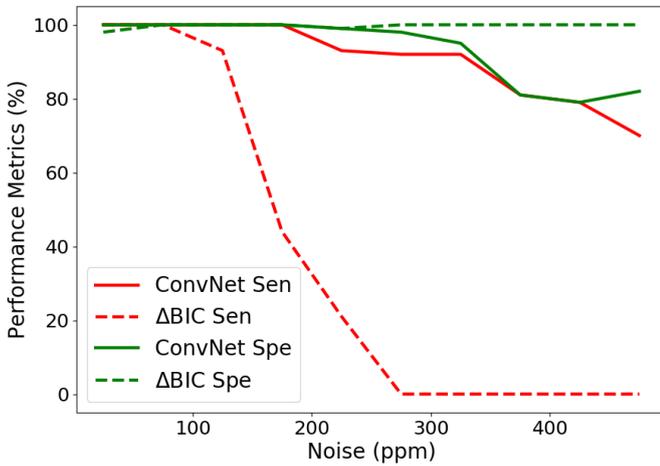


Fig. 5. Performance as a function of noise level between 25 ppm and 475 ppm. The exomoon radius is $1 R_{\oplus}$. Data set 3 was used.

to $\text{Acc} = 75\%$ for the highest noise level. The sensitivity is higher for the CNN than ΔBIC for two reasons. First, the total variation (Beck & Teboulle 2009) enforces spatial constraints, which helps the learning process to define spatial continuity in the signal and distinguish between a moon-like signal and noise. Second, adding Gaussian noise with dropout also helps to detect changes in the signal (see, e.g., Srivastava et al. 2014).

On the other hand, for high noise levels, the difference in the maximum likelihood between the one-moon model and the no-moon model is too small to counteract the penalty given by the difference in the number of parameters between the two models.

4.5. Effect of training size

Figure 6 presents the performance metrics as functions of the number of samples in the training set from 300 000 to 700 000. The number of samples for the testing set was kept constant at 315 720 samples. As expected, the CNNs are more robust when larger training sets were used. The results show that training sets with 650 000 samples achieve only little improvement. For the largest three training sets shown in Fig. 6, the CNN performance appears to have converged.

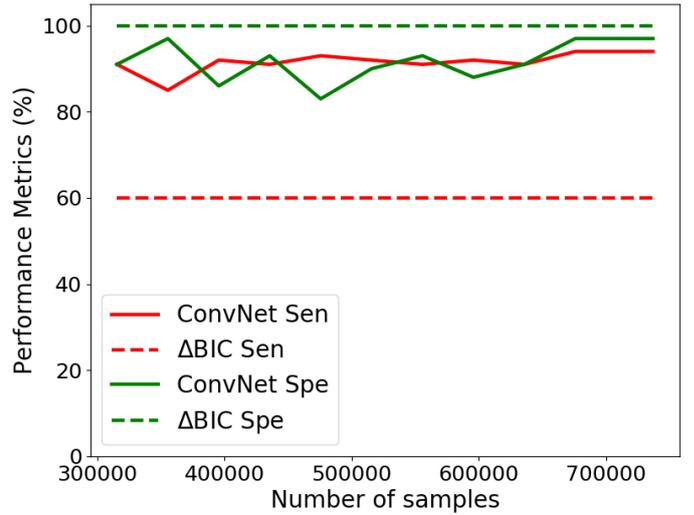


Fig. 6. Results as a function of training size. The training set includes all noise levels and exomoon sizes. Data set 1 was used.

Table 6. Performance comparison between the proposed CNN and other 1D CNN architectures.

Method	Sen (%)	Spe (%)	Acc (%)	Computation time (s)
ΔBIC	60	100	80	$>2 \times 10^{10}$
Vgg16	97	86	86	164
AlexNet	0	100	50	90
ResNet	44	100	75	170
DenseNet	44	100	75	180
Proposed	94	97	96	158

Notes. The last column gives the computing time required to process the test data. Data set 1 was used.

4.6. Comparison with the ΔBIC method and other CNN architectures

Table 6 shows the comparison between the method we propose here, the classic approach for detecting moon-like signals (ΔBIC), and other well-known CNN architectures (Vgg16 1D, Liu & Deng 2015; AlexNet 1D, Krizhevsky et al. 2012; ResNet 1D, He et al. 2016; and DenseNet 1D, Huang et al. 2017). The difference between the approach we propose and the other CNN approaches is in the size of the kernel filters, the number of filter units, the number of convolution layers, and the order of convolution layers. The results of our proposed CNN with total variation loss outperforms all other approaches. The classical approach (ΔBIC) performs more poorly than all CNNs.

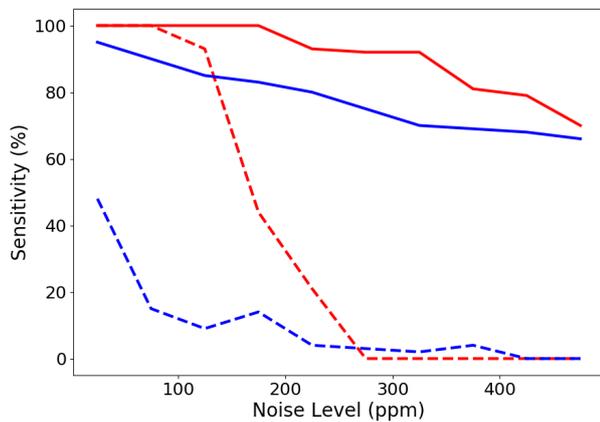
The ΔBIC method requires around 12–18 h per light curve on one CPU with a memory of 32 GB per core, which requires above 2×10^{10} s to analyze 315 720 light curves (test data). On the other hand, the CNN requires less than one second per light curve on one GPU with 32 GB per core, which corresponds to only about 200 seconds for the entire test data set.

4.7. Training the CNN using a more general data set

The synthetic data from the previous sections were all generated using only three free parameters (exomoon radius, noise level, and exomoon phase), while the other parameters of the star and planet were kept constant. In order to test how the CNN method

Table 7. Stellar, planetary, and exomoon parameters for the transit light-curve simulations with 14 free parameters.

Parameter	Range
Planet-star radius ratio, r_p	0.065–0.126
Barycenter semimajor axis, a_b	$49.5R_\star$ – $281.7R_\star$
Barycenter impact parameter, b	–1 to 1
Barycenter initial orbital phase, φ_b	0
Barycenter orbital period, P_b	200–300 days
Stellar limb-darkening parameters	$q_1 = 0.4$ – 0.6 $q_2 = 0.4$ – 0.6
Moon-star radius ratio, r_m	0.0065–0.0504
Moon semimajor axis, a_m	$10^{-5}R_\star$ – $0.005R_\star$
Moon initial orbital phase, φ_m	full range
Moon orbital period, P_m	0.5–10 days
Moon-planet mass ratio, f_m	0.001–0.2
Moon orbital inclination	full range
Moon argument of ascending node	full range
Noise level	50–500 ppm

**Fig. 7.** Sensitivity as a function of noise level using input data generated with 3 free parameters (exomoon radii, exomoon orbital phases, and noise levels, see Table 2; red curves) and with 14 free parameters (see Table 7; blue curves). The two methods are compared at a fixed exomoon radius of $1 R_\oplus$: ConvNet (solid curves) vs. Δ BIC (dashed curves). The sensitivity is plotted in bins of 50 ppm noise levels.

behaves when it is trained on a more general data set, we generated new synthetic light curves for which we varied all 14 free parameters (Table 7). This new data set includes 200 000 light curves.

Figure 7 shows the sensitivity of both Δ BIC and the CNN versus noise level at a fixed exomoon radius of $1 R_\oplus$. When the CNN is trained using the dataset with 14 free parameters, the performance of the CNN decreases by up to 20% compared to its performance when the more restricted training data set is used. It is important to note that the CNN still performs well at high noise levels. For noise levels above about 200 ppm, the performance of the Δ BIC method is far worse than the CNN, independently of the data set it is applied to.

5. Conclusion

We proposed a regularized 1D CNN architecture to detect exomoon candidates in transit light curves. The regularized loss function brings a significant improvement over traditional loss

functions in terms of convergence, sensitivity, and specificity. Furthermore, the method we propose performs relatively well with high levels of noise and small moon sizes.

We compared the network results with the Δ BIC method described by Rodenbeck et al. (2018). The Δ BIC method is more conservative: it has a higher true-negative rate, but also a much lower true-positive rate. For low noise levels (100 ppm), both methods are able to reliably detect exomoons larger than $1 R_\oplus$ in our simulation setup. The CNN shows its strength at higher noise levels. For the data set with only 3 free parameters, Sen and Spe stay above 90% up to a noise level of 325 ppm and remain around 80% for even higher noise levels. For the data set with 14 free parameters, Sen remains above 65% at 500 ppm. The Δ BIC method is completely insensitive to the presence of exomoon at noise levels of 275 ppm and higher. The CNN is more effective in predicting a moon-like signal, and it is faster. It is a promising technique for applications to real astronomical data sets.

Acknowledgements. We thank Chris Hanson for constructive comments. This work was supported in part by NYUAD Institute grant G1502 “Center for Space Science”. RA acknowledges support from the NYUAD Kawader Research Program. Computational resources were provided by the NYUAD Institute through Muataz Al Barwani and the HPC Center. The simulated light curves were produced at the DLR-supported PLATO Data Center at the MPI for Solar System Research. Author Contributions: LG proposed the research idea. RA proposed the Machine Learning algorithm and analyzed the data and the results. KR prepared the simulated data and performed the BIC analysis. KR is a member of the International Max Planck Research School for Solar System Science at the University of Göttingen. All authors contributed to the final manuscript.

References

- Beck, A., & Teboulle, M. 2009, *IEEE Trans. Image Proc.*, **18**, 2419
- Claret, A., & Bloemen, S. 2011, *A&A*, **529**, A75
- Hand, D. J., Smyth, P., & Mannila, H. 2001, *Principles of Data Mining* (MIT Press)
- He, K., Zhang, X., Ren, S., & Sun, J. 2016, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770
- Heller, R. 2014, *ApJ*, **787**, 14
- Hippke, M. 2015, *ApJ*, **806**, 51
- Houhou, N., Bresson, X., Szlam, A., Chan, T. F., & Thiran, J.-P. 2009, in *Scale Space and Variational Methods in Computer Vision*, eds. X.-C. Tai, K. Mørken, M. Lysaker, & K.-A. Lie, 112
- Huang, G., Liu, Z., van der Maaten, L., & Weinberger, K. Q. 2017, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 4700
- Ioffe, S., & Szegedy, C. 2015, Proceedings of the 32nd International Conference on Machine Learning (Lille, France: PMLR), eds. F. Bach, & D. Blei, *Proc. Mach. Learn. Res.*, **37**, 448
- Janocha, K., & Czarnecki, W. M. 2017, *Theor. Found Mach. Learn.*, **25**, 49
- Javanmardi, M., Sajjadi, M., Liu, T., & Tasdizen, T. 2016, *IEEE Conf. Comput. Vision Pattern Recogn.*
- Kingma, D. P., & Ba, J. 2015, in 3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings, eds. Y. Bengio, & Y. LeCun
- Kipping, D. M. 2009, *MNRAS*, **392**, 181
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. 2012, *Adv. Neural Inf. Proc. Syst.*, **25**, 1097
- Liu, S., & Deng, W. 2015, *3rd IAPR Asian Conference on Pattern Recognition*, 730
- Maas, A. L., Hannun, A. Y., & Ng, A. Y. 2013, *ICML 2013 Workshop on Deep Learning for Audio, Speech and Language Processing*, Atlanta, GA, USA, June 16, 2013
- Mahendran, A., & Vedaldi, A. 2015, *Int. J. Comput. Vision*, **120**, 233
- Pedamonti, D. 2018, *IEEE Conf. Comput. Vision Pattern Recognit.*
- Rodenbeck, K., Heller, R., Hippke, M., & Gizon, L. 2018, *A&A*, **617**, A49
- Schwarz, G. 1978, *Ann. Stat.*, **6**, 461
- Simonoff, J. S. 1998, *Smoothing Methods in Statistics* (Springer)
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. 2014, *J. Mach. Learn. Res.*, **15**, 1929
- Szabó, R., Szabó, G. M., Dály, G., et al. 2013, *A&A*, **553**, A17
- Teachey, A., & Kipping, D. M. 2018, *Sci. Adv.*, **4**
- Teachey, A., Kipping, D. M., & Schmitt, A. R. 2018, *AJ*, **155**, 36

Appendix A: Network architecture

The network consists of the following: convolution blocks (convolution, batch normalization, maximum pooling, leaky ReLU activation and Gaussian dropout), fully-connected block and classification block.

First convolution layer ($\ell = 0$). The input of the first convolution layer, $\ell = 0$, is a vector X_m of dimension N . The index m of the light curve is fixed throughout this section, therefore we dropped it to simplify the notation. To construct feature maps, the input data were convolved 64 times with different kernels $F_{k,0}$ of width $n_f^0 = 3$ and stride $s_f = 1$. The kernel values are to be determined later by training the network. The stride is the distance used to shift the convolution kernel. We obtained 64 feature maps of length $N_f^0 = N$ (with zero padding), see, for example, Krizhevsky et al. (2012),

$$Y_{\text{conv}}^{k,0}(i) = \sum_{j=0}^{n_f^0-1} X(i-j)F_{k,0}(j), \quad 0 \leq k < K_0 = 64. \quad (\text{A.1})$$

Batch normalization. The convolution operator is followed by a batch normalization function. Batch normalization is applied to reduce internal covariance shift (Ioffe et al. 2015). It is simply a linear transformation of the data with scaling λ_0 and a shift β_0 ,

$$Y_{\text{norm}}^{k,0}(i) = \lambda_0 \frac{Y_{\text{conv}}^{k,0}(i) - \mu_0}{\sqrt{v_0 + \epsilon}} + \beta_0, \quad (\text{A.2})$$

where μ_0 and v_0 are the average and variance of the input data for all ks at fixed ℓ . The initial values of the parameters were $\lambda_0 = 1$ and $\beta_0 = 0$. The constant $\epsilon = 0.002$ was added to avoid division by very low values.

Leaky ReLU activation. The batch normalization is followed by applying a leaky rectified linear unit (ReLU) activation function as follows (Maas et al. 2013):

$$Y_{\text{lactiv}}^{k,0}(i) = \text{LReLU}\left(Y_{\text{norm}}^{k,0}(i); w_k, b_k\right) = \begin{cases} Y_{\text{norm}}^{k,0}(i)w_k + b_k & \text{if } Y_{\text{norm}}^{k,0}(i) \geq 0, \\ \rho Y_{\text{norm}}^{k,0}(i)w_k + b_k & \text{else.} \end{cases} \quad (\text{A.3})$$

Here w and b are values to be determined. Many activation functions are possible, such as tanh, sigmoid, hyperbolic, or rectified functions. The leakly ReLU function allows for a small gradient ρ when the units are not active. Here, we set $\rho = 0.02$, and refer to Maas et al. (2013) and Pedamonti (2018) for an in-depth discussion.

Maximum pooling. A maximum-pooling operator performs a spatial subsampling in time by taking the maximum value over a pooling window of length n_p every s_p points (with $n_p = s_p$),

$$Y_{\text{mpool}}^{k,0}(i) = \max\{Y_{\text{lactiv}}^{k,0}(j)\}_{i s_p \leq j \leq i s_p + n_p - 1}, \quad (\text{A.4})$$

for $i = 0, 1, \dots, N_f^0/n_p - 1$.

Gaussian dropout. In order to provide some regularization, a Gaussian dropout procedure is applied after the maximum pooling (Srivastava et al. 2014). It involves multiplying weights by a Gaussian random variable with mean 1 and standard deviation $\sigma = 0.65$. Symbolically, we write

$$Y_{\text{dout}}^{k,\ell} = \text{Gaussian_dropout}\left(Y_{\text{mpool}}^{k,\ell}\right). \quad (\text{A.5})$$

We found through experimentation that a Gaussian dropout gives better results than a standard dropout.

Next convolution layers. Another four convolution layers ($\ell = 1, 2, 3$, and 4) follow the first convolution layer ($\ell = 0$). For each layer $\ell \geq 1$, we computed the feature maps,

$$Y_{\text{conv}}^{k,\ell}(i) = \sum_{j=0}^{n_f^{\ell-1}} Y_{\text{dout}}^{k,\ell-1}(i-j)F^{k,\ell}(j), \quad 0 \leq k < K_\ell, \quad (\text{A.6})$$

followed by batch normalization, leaky ReLU activation, maximum pooling, and dropout, as defined above. The number of feature maps K_ℓ is 128 for the intermediate convolution layers $\ell = 1, 2, 3$ and $K_4 = 64$ for the last layer.

Fully connected layer. The output of the last convolution layer $\ell = 4$ is a matrix of dimension 64 (number of feature maps) times 102 (number of time samples). This layer is followed by a fully connected layer $Y(p)$ activated by a leaky ReLU (see Eq. (A.3)),

$$Y(p) = \sum_{k=0}^{63} \sum_{i=0}^{101} \text{LReLU}\left(Y_{\text{dout}}^{k,4}(i); w_{k,i}(p), b(p)\right), \quad (\text{A.7})$$

for $p \leq 0 \leq N_p - 1$, where $N_p = 64 \times 102 = 6528$ is the number of neurons; this is linearly converted to 2048. The fully connected layer is the final vector output before the data are classified.

Output Layer. Sigmoid activation and classification. This layer is activated by a sigmoid function, which produces a probability in the range between 0 and 1, as follows:

$$\hat{y} = (1 + e^{-y})^{-1}, \quad (\text{A.8})$$

with

$$y = \sum_{p=0}^{N_p-1} w_p Y(p) + b. \quad (\text{A.9})$$

The final classification of the data is

$$C(X) = \begin{cases} 1 & \text{if } \hat{y} \geq 0.5, \\ 0 & \text{else.} \end{cases} \quad (\text{A.10})$$

Appendix B: Optimization

We used a batch approach, whereby the input samples are passed forward and backward through the network in chunks of 32 samples (in each epoch). We refer to Kingma & Ba (2015) for additional details about this algorithm (AdaMax). Denoting Θ_{q-1} the parameters of the network at iteration $q - 1$, the updated values at iteration q are given by

$$\Theta_q = \Theta_{q-1} - \left(\frac{\alpha}{1 - \beta_1}\right) \frac{m_q}{\mu_q}, \quad (\text{B.1})$$

with

$$m_q = \beta_1 m_{q-1} + (1 - \beta_1) \nabla L(\Theta_{q-1}), \quad (\text{B.2})$$

and

$$\mu_q = \max(\beta_2 \mu_{q-1}, \|\nabla L(\Theta_{q-1})\|). \quad (\text{B.3})$$

The quantity α is the step size. $\eta = \alpha/(1 - \beta_1)$ is the learning rate, and the parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$ were fixed.

Initially, we set $\alpha = 0.002$. At $q = 0$, we initialized the weights of the network in each layer with a random number drawn from a zero-mean Gaussian distribution with standard deviation 0.1. The biases were initially set to zero.

A stopping criterion was used to reduce overfitting of the network. Training stopped when the validation error did not decrease after 20 epochs. If the loss function did not decrease, the parameter α was multiplied by a factor of 0.1 after the first 10 epochs, and the operation was repeated every 10 epochs.

Figure B.1 compares the evolution of the loss function versus the number of epochs used to train the network. The regularized loss function performs the best.

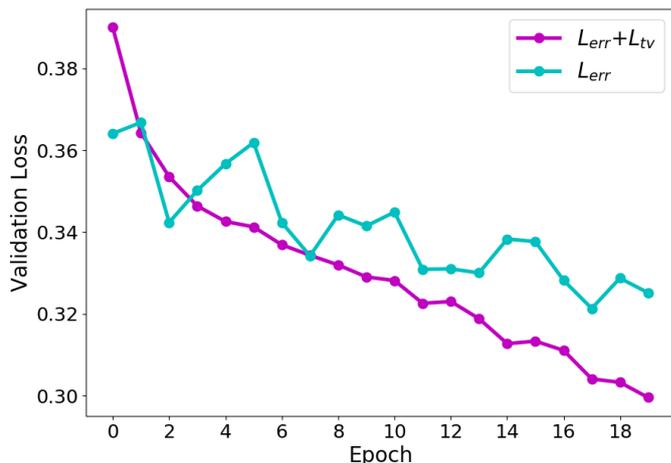


Fig. B.1. Validation loss as a function of the number of epochs used to train the network. An epoch consists of 32 light curves. The regularized loss function $L_{err} + L_{tv}$ decreases faster than L_{err} alone.