



# Shallow Transits—Deep Learning. II. Identify Individual Exoplanetary Transits in Red Noise using Deep Learning

Elad Dvash<sup>1</sup>, Yam Peleg<sup>1</sup>, Shay Zucker<sup>1</sup> , and Raja Giryes<sup>2</sup>

<sup>1</sup> Porter School of the Environment and Earth Sciences, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, 6997801, Israel; [shayz@tauex.tau.ac.il](mailto:shayz@tauex.tau.ac.il)

<sup>2</sup> School of Electrical Engineering, Iby and Aladar Fleischman Faculty of Engineering, Tel Aviv University, Tel Aviv, 6997801, Israel

Received 2019 November 26; revised 2022 March 14; accepted 2022 March 15; published 2022 April 27

## Abstract

In a previous paper, we introduced a deep learning neural network that should be able to detect the existence of very shallow periodic planetary transits in the presence of red noise. The network in that feasibility study would not provide any further details about the detected transits. The current paper completes this missing part. We present a neural network that tags samples that were obtained during transits. This is essentially similar to the task of identifying the semantic context of each pixel in an image—an important task in computer vision, called “semantic segmentation,” which is often performed by deep neural networks. The neural network we present makes use of novel deep learning concepts such as U-Nets, Generative Adversarial Networks, and adversarial loss. The resulting segmentation should allow further studies of the light curves that are tagged as containing transits. This approach toward the detection and study of very shallow transits is bound to play a significant role in future space-based transit surveys such as PLATO, which are specifically aimed to detect those extremely difficult cases of long-period shallow transits. Our segmentation network also adds to the growing toolbox of deep learning approaches that are being increasingly used in the study of exoplanets; but, so far mainly for vetting transits, rather than their initial detection.

*Unified Astronomy Thesaurus concepts:* [Exoplanet detection methods \(489\)](#); [Transit photometry \(1709\)](#); [Neural networks \(1933\)](#)

## 1. Introduction

In a previous paper (Zucker & Giryes 2018; hereafter Paper I), we demonstrated a new approach to detect the presence of exoplanetary transits in simulated data mimicking data that can be obtained by high-cadence space telescopes. The demonstration was performed on simulated data of a fictitious telescope, but the approach should be applicable to real-life missions like CoRoT (Deleuil et al. 2010), Kepler (Borucki et al. 2010), TESS (Ricker et al. 2015), and in the future PLATO (Rauer et al. 2016). The new approach we have introduced aimed at overcoming the problem of “red noise,” usually attributed mainly to stellar activity, which constituted a major hurdle to traditional transit detection techniques, such as the BLS (Kovács et al. 2002). Our suggested approach was based on the rapidly evolving new discipline of deep learning. In Paper I we demonstrated how this technique managed to outperform the BLS (preceded by a high-pass filter), in identifying light curves that contained exoplanetary transits, contaminated by red noise in addition to photon (Poisson) white noise.

It is important to note that Paper I focused on the task of detecting the presence of transits in the light curves, and not validating or vetting them as exoplanetary signals. The aim was to detect those transit events that might evade detection by traditional detection approaches like the BLS. In that respect it differed from other efforts in the field (e.g., Ansdell et al. 2018; Shallue & Vanderburg 2018; Dattilo et al. 2019; Osborn et al. 2019; Yu et al. 2019).

As successful as it may be, the detection mechanism we had introduced in Paper I lacked one crucial ingredient: it could not provide any information as to the details of the detected transits. The information it provided was binary: whether the light curve contained transits or not. Our aim in the current work is to present a deep learning neural network that will also identify the individual transits in the light curve, thus enabling further research, such as vetting the transit candidates, characterizing the transit properties, detecting transit timing variations (TTVs), looking for additional transiting planets, etc.

Deep learning is a class of algorithms and heuristics meant to train highly nonlinear parametric functions. The nonlinear functions, mostly known as neural networks, are essentially concatenations of layers of basic units, each comprising a linear operation followed by a simple nonlinearity. The nonlinearity is commonly realized by element-wise activation functions such as the sigmoid, hyperbolic tangent, or the rectified linear unit (Nair & Hinton 2010). Their combination eventually results in intricate highly nonlinear functionality.

During the training, the parameters of each layer are trained so as to minimize an error function calculated in relation to the previous layer. The training is often done using stochastic gradient descent (Rumelhart et al. 1986). This approach often captures strong nonlinear relationships, leading to unprecedentedly successful results across many fields (e.g., LeCun et al. 2015; Schmidhuber 2015; Goodfellow et al. 2016).

The task of identifying samples that are included in individual transits is essentially equivalent to the task of “semantic segmentation” in computer vision. The goal of the segmentation task is usually to simplify and change the representation of an image into something that is more meaningful and easier to analyze. Image segmentation is essentially the partitioning of a digital image into multiple



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

segments, usually corresponding to objects and boundaries (lines, curves, etc.) in the image. Thus, image segmentation can be described as the process of assigning a label to every pixel in an image such that pixels with the same label share certain characteristics, or simply belong to the same context. The equivalence to the task of identifying the transits in a light curve is obvious: we assign a label to each sample in the light curve such that all the samples within transits get the same label. Since much progress has been achieved in performing image segmentation using deep learning, it is only natural to apply it here as well.

Most of the aforementioned studies, which aimed at detecting, vetting, and identifying transits, made use of convolutional neural networks. In the current work, we use more tools from the toolkit of neural networks to perform segmentation. In particular, we use U-Nets (Ronneberger et al. 2015) to perform the segmentation and identify the times when a transit occurs within a given light curve signal, and an adversarial loss to force the network to output only realistic segmentation.

In the next section we introduce the neural network concepts that we employed in our work. In Section 3 we present the way we implemented those concepts in our neural network. Section 4 describes the simulated data set used for our demonstration, and the procedure we used to train the network is detailed in Section 5. Section 6 demonstrates the performance of the neural network, and in Section 7 we conclude and discuss the possible future implementations of the approach.

## 2. Neural Networks

The approach we suggest for solving the problem of identifying individual transits makes use of several variants of deep learning neural networks: convolutional neural networks (CNNs), ResNets, U-Nets, and generative adversarial networks (GANs). In the next paragraphs we briefly introduce and explain these concepts, as well as other concepts we employ.

### 2.1. Convolutional Neural Networks

In CNNs, convolutions constitute the linear part of the layers (LeCun et al. 1998). CNNs are widely used to analyze images or periodic signals due to their shift-invariance property. CNNs are usually built by stacking convolution operators in layers, each followed by a nonlinearity (“activation function”). Usually, the stack of convolutions is followed by a fully connected layer, represented by a simple linear function (matrix multiplication), and followed by an activation function. These networks are known to be very powerful when applied to signal classification tasks (see Paper I). The layers are usually “contracting,” in the sense that they perform successive downsampling of the signal, resulting in an increasingly compact representation of the information.

### 2.2. Residual Networks (ResNets)

An essential step in training neural networks is back propagating the gradient of the loss function through the layers. A notorious problem in training very deep networks is the problem of “vanishing gradients”: during the back-propagation of the gradients, repeated multiplications cause the gradients to become too small for effective learning. As a result, as networks grow deeper, the performance plateaus and

might even start to degrade. A standard technique to avoid this problem uses residual connections (also known as skip connections): retaining the original output of an earlier layer and adding it to the results of following layers as a “bypass” (He et al. 2016). This helps to mitigate the vanishing gradient problem by causing the gradient from the earlier layer to flow through the bypass and skip multiplication steps.

### 2.3. Fully Convolutional Networks and U-Nets

Building upon the concept of a CNN, a more elaborate deep learning architecture has emerged—the fully convolutional network (FCN), which is very popular for image segmentation (Long et al. 2015). A popular FCN structure is the U-Net (named after the U-shape of the network), which has been initially developed for biomedical images but has become widely used in many domains (Ronneberger et al. 2015). Essentially, it is a CNN that is composed of two parts, the “encoder” and the “decoder.” The encoder is a contracting CNN, which produces a compact representation of the input signal. The decoder, which is appended to the encoder, comprises mirrored layers (with respect to the encoder), in the sense that each convolution in the encoder is mirrored by a corresponding deconvolution (transposed convolution) layer. As a consequence, this expansive path is more or less symmetric to the contracting path, yielding a U-shaped architecture. It has been found that these networks perform better with the following improvement: depth-wise concatenation of the output of each encoding layer to the corresponding decoding layer in the mirrored architecture (Ronneberger et al. 2015). By design, the original U-Net takes two-dimensional, single-channel (grayscale) images as inputs. The current study deals with light curves, which are one-dimensional (1D) time series. We have therefore restructured the U-Net design to take 1D time series as inputs by using 1D convolution layers.

### 2.4. Dice Loss

The objective of the training of a neural network is the minimization of a prescribed *loss function*. The loss function represents the task one wishes the network to perform, and its goal is to provide a metric that measures the performance of the network for the given task. While in the problem of classification (e.g., Paper I) the loss function commonly used is the logarithmic loss (also known as the cross-entropy loss), in this work we chose to apply a variant of the *Dice Loss*, which is useful for segmentation problems.

Historically, the Dice coefficient was inspired by a set-theoretic concept introduced independently by Dice (1945) and Sørensen (1948) in ecological contexts in order to quantify similarity of sets. In the set-theoretic context, the Dice coefficient of the two sets  $X$  and  $Y$  is defined by:

$$d(X, Y) = \frac{2|X \cap Y|}{|X| + |Y|} \quad (1)$$

where  $|\cdot|$  denotes the number of elements in each set. The concept of set membership is generalized to binary sequences, and the Dice coefficient for two binary sequences  $\{y_i\}$  and  $\{p_i\}$

can now be written as:

$$d(y, p) = \frac{2\sum_i p_i y_i}{\sum_i p_i^2 + \sum_i y_i^2}. \quad (2)$$

Milletary et al. (2016) were the first to apply the Dice coefficient to image segmentation. In this context, it is essentially a measure of the overlap between the segmentation image that the network produces and the ground-truth segmentation sequence. In our context, for a given light curve, let us denote the ground truth by a binary sequence  $\{y_i\}$ , where each sample in transit is assigned the value 1 whereas all the rest are assigned 0. Let  $\{p_i\}$  denote the output of our segmentation network (the “prediction,” in machine learning jargon), which is the probability (a value between 0 and 1) that sample  $i$  is within a transit. We wish this probability to be 1 for samples during transit and 0 otherwise. Then the Dice coefficient for this light curve is given by:

$$l = d(y, p) + d(1 - y, 1 - p), \quad (3)$$

where  $d(y, p)$  measures the performance during transit segments and  $d(1 - y, 1 - p)$  during out-of-transit segments.

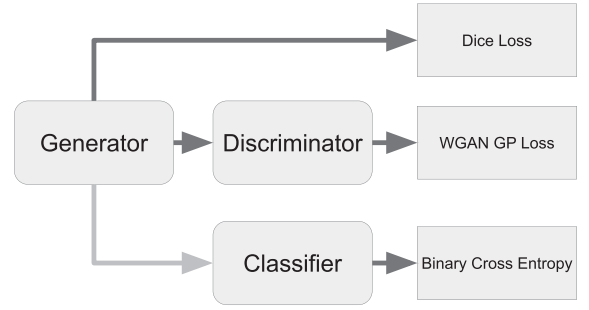
### 2.5. Adversarial Loss

Our aim in the current project is to label samples occurring in transits. Naively, under the assumption that transits are strictly periodic (neglecting TTV and multiple transiting planets), it should have been very easy to judge whether the results of the segmentation are realistic. However, we wished to leave our mechanism agnostic of the strict periodicity of the signals to allow, in future developments, the detection of multiplanetary signals, or signals with significant TTV. Thus, it is quite difficult to define a metric to measure how realistic is the resulting segmentation. Therefore, in training our neural network we make use of a novel concept in deep learning which is the concept of a GAN.

Maximizing the Dice loss (Equation (3)) forces the neural network to output a segmentation that is similar to the ground-truth one. However, in this loss function there is no preference for the transit signal to be necessarily periodic. Thus, the network might produce predictions that minimize the Dice loss but do not look “authentic,” i.e., similar to real transits. An experienced exoplanet astronomer that would examine such an output would immediately be able to exclude an unauthentic segmentation. It is therefore required to add some kind of a penalizing mechanism to the network during training, so as to exclude those false predictions.

This problem is not unique to our setup only, but is common in the training of neural networks. There is a trade-off between minimizing the distortion (the Dice loss in our case) and the naturalness of the reconstructed signal (see, for example, the analysis for the case of superresolution by Blau & Michaeli 2018). Thus, one may add a loss term in the training of the neural network that pushes the output distribution to resemble the true data distribution. GAN is a very popular strategy for achieving this goal.

A GAN comprises two neural networks, where one network (the “generator”) generates “candidate” signals and the other (the “discriminator”) evaluates them and discriminates between actual signals and ones produced by the generator. The training objective of the generator is to challenge the discriminator and increase its error rate (i.e., “fool” it by producing novel synthesized instances that appear to be genuine and cannot be



**Figure 1.** A schematic illustration of the overall network structure. Note the gray arrow between the generator and classifier signifying no gradients flow through those connections.

distinguished from real data). The discriminator discriminates between genuine instances and artificial candidates produced by the generator (e.g., Goodfellow et al. 2014).

Usually, GANs are used to generate purely new signals from some initial distribution. However, in our implementation, we use the generator to tag the samples that occur during transits, essentially producing a new sequence. The discriminator examines the resulting sequence and evaluates how realistic it is as a sequence of transit events.

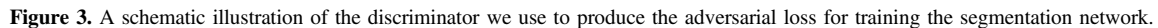
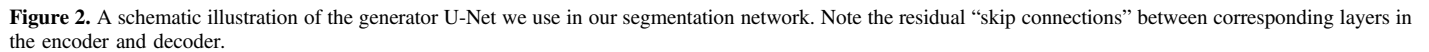
GANs are known to suffer from training instability. In particular, a known problem in their training is “mode collapse,” where the generated examples represent only a small fraction of the real distribution (e.g., the generator might generate always the same real-looking image). As a solution, a variant of GAN called Wasserstein GAN (WGAN) has been proposed in which the loss of the discriminator is set to be the Wasserstein distance (a measure of the distance between two probability distributions, also known as the Earth Mover’s Distance) leading to a more stable training (Arjovsky et al. 2017).

However, WGAN suffers from another problem, which is exploding of the gradient norm. As a solution to this problem, yet another variant of GAN has been proposed—the WGAN gradient penalty (WGAN-GP), which penalizes the norm of the gradient of the discriminator network (Gulrajani et al. 2017). This method is known to perform even better, as it enables stable training of a wide variety of GAN architectures with almost no hyperparameter tuning.

In our case, we may train a discriminator that distinguishes between the output of the segmentation network (which acts as the generator) and the real ground-truth segmentation sequences. This use of the GAN framework to improve training of a given network (e.g., our segmentation network) is known as adding an adversarial loss, since the discriminator in this case is used as an “additional loss” in the training of the segmentation network. Specifically, in addition to maximizing the Dice coefficient, our segmentation network also aims at “fooling” the discriminator, which makes its output more similar to the ground-truth data.

### 3. Current Implementation

For the task of transit segmentation, we use the U-Net encoder-decoder architecture followed by the WGAN discriminator. In this setup, one may consider the U-Net as the generator of the WGAN network. Our model facilitates joint detection and segmentation with one architecture. Figure 1 presents the overall structure of this architecture. Notice that it

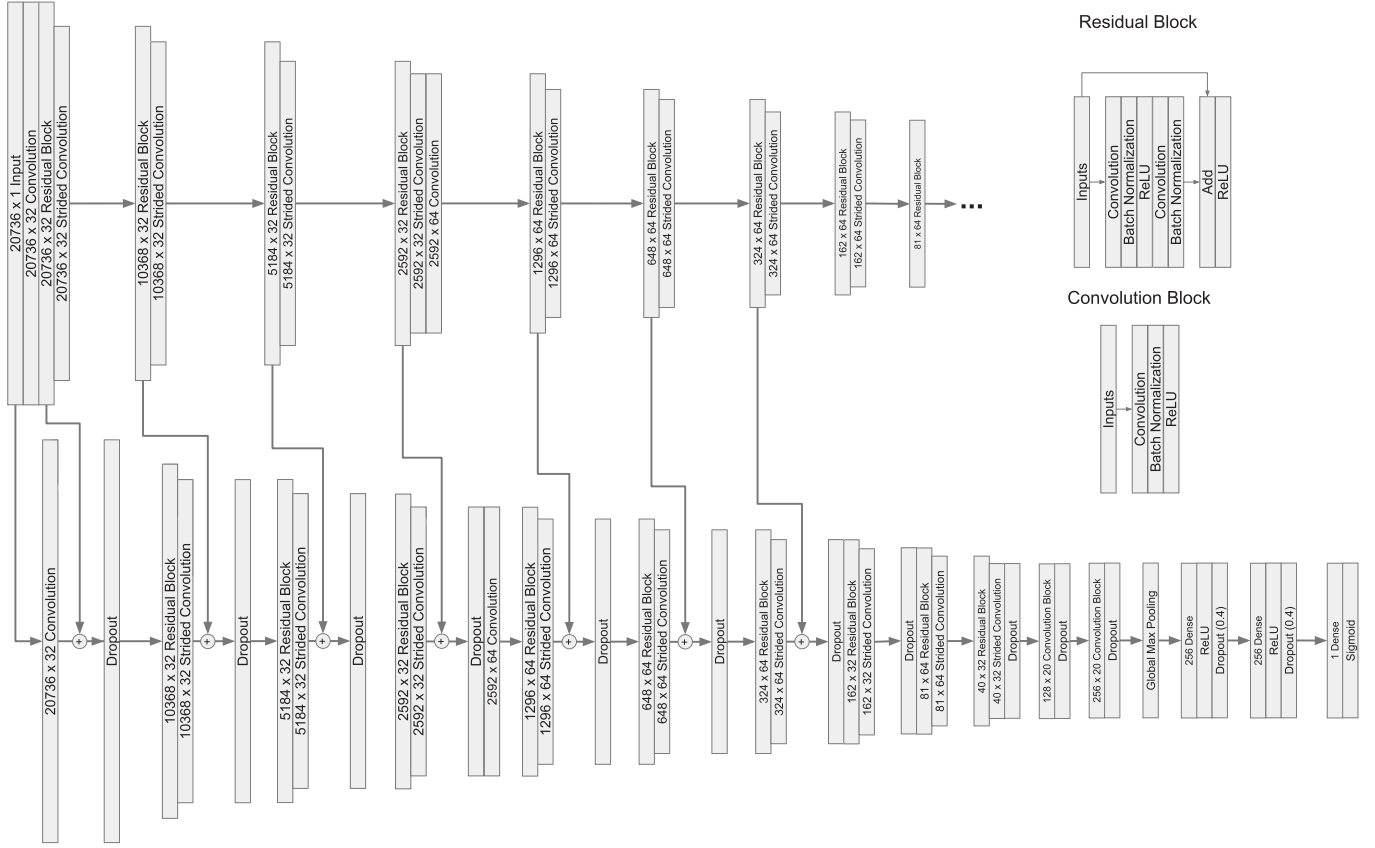


The training process is split into two parts. First, we train the generator and discriminator only on time series that contain transits, as a regular WGAN-GP, using the Dice loss combined with the adversarial loss, weighted 0.75 and 0.25, respectively, when training the generator. In the second part of each training iteration, we freeze the weights of the generator and discriminator and train the classifier using binary cross entropy on light curves, where only some of them contain transit signals. The fine technical details of the various architectures we used can be found in our code, which is publicly available on GitHub<sup>3</sup> and archived in Zenodo (Dvash et al. 2022).

The magnitude directly affects the uncorrelated photon noise, in a way that should depend on the characteristics of the observational apparatus. In order to estimate the white noise

4





**Figure 4.** A schematic illustration of the classifier network, which uses the residual connections from the generator U-Net.

component in a two-minute cadence TESS light curve, we approximated the curve in Figure 4 of Ricker et al. (2015) by the following relation (assuming the brightness is quantified in the  $I_C$  band):

$$A_w = \sqrt{108000 + 8670 e^{0.94(I_C - 5)}} \mu\text{mag}. \quad (4)$$

Following Paper I, we have used a Gaussian process (GP) to simulate the noise, with a kernel comprising a squared-exponential component and a quasiperiodic one (e.g., Aigrain et al. 2016). Combined with the white noise component, we got the following expression for the kernel of the noise GP:

$$(t_i, t_j) = A_s^2 \exp \left[ -\left( \frac{t_i - t_j}{\lambda_s} \right)^2 \right] + A_q^2 \exp \left[ -\frac{1}{2} \sin^2 \left( \frac{\pi(t_i - t_j)}{T_q} \right) - \left( \frac{t_i - t_j}{\lambda_q} \right)^2 \right] + A_w^2 \delta(t_i - t_j). \quad (5)$$

We assumed that the details of the red noise were related to stellar properties and not to the observational apparatus, and therefore we used the same distributions for the non-white-noise components as we had previously used in Paper I. Table 1 summarizes the various hyperparameters of the GP that we have used. Note that unlike the case in Paper I we have not added artificially any outlier samples to the noise.

Unlike in Paper I, the ability to perform the segmentation might be affected by the detailed shape of the transits. Therefore, we could no longer settle for a simple trapezoid transit model. Instead we chose to use the publicly available

**Table 1**  
GP Kernel Hyperparameter Ranges

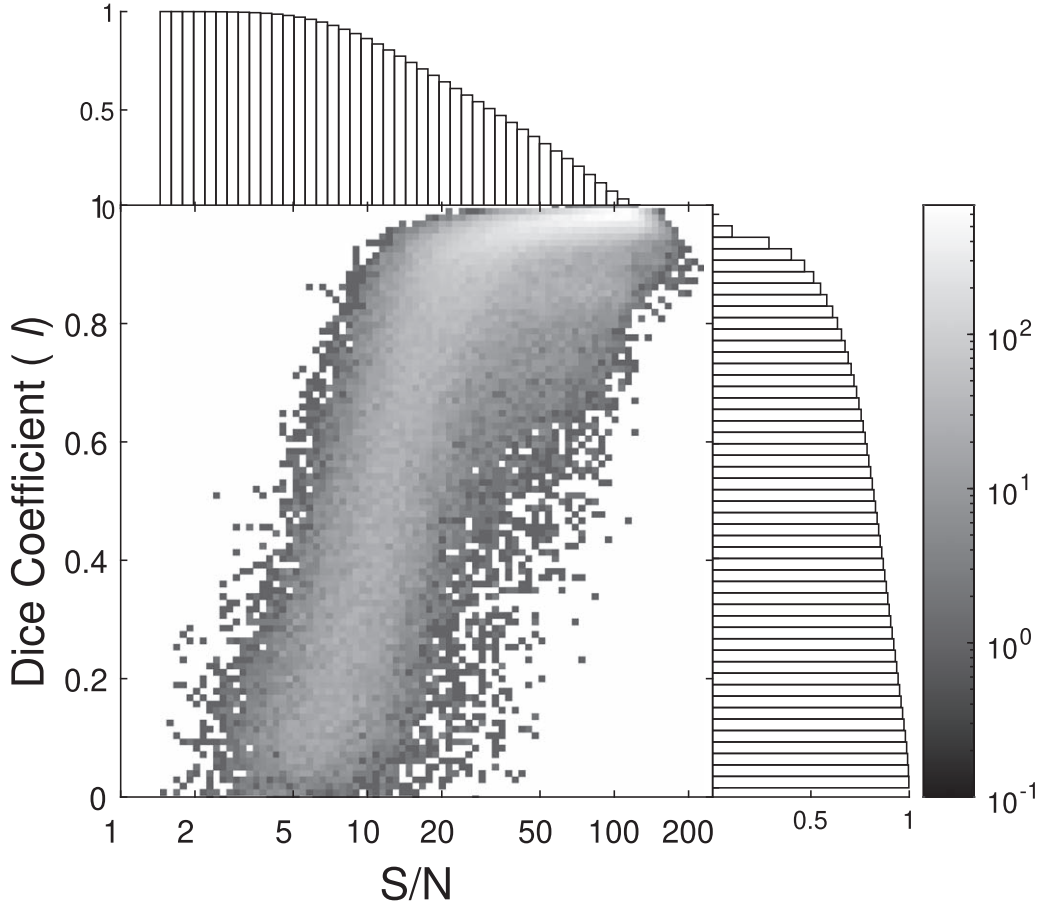
Hyperparameter	Minimum Value	Maximum Value
$A_s$	5 $\mu\text{mag}$	125 $\mu\text{mag}$
$A_q$	50 $\mu\text{mag}$	125 $\mu\text{mag}$
$\lambda_s$	1 minutes	10 hr
$T_q$	10 hr	500 hr
$\lambda_q$	1000 minutes	500 hr

**Table 2**  
Distributions of Simulated Transit Parameters

Parameter	Distribution	Minimum Value	Maximum Value
$P$	Log-uniform	2 days	9 days
$R_p/R_*$	Log-uniform	1/100	1/30
$b$	Uniform	0	1
$c_1$	Uniform	0.5	0.7

code BATMAN which is capable of simulating transits quickly and accurately in a wide range of parameters and with various options to simulate the limb darkening (Kreidberg 2015). We chose to use a linear limb-darkening model (e.g., Howarth 2011), with a single parameter  $c_1$ .

We drew a sample of stellar masses using a Salpeter initial mass function between  $0.3 M_\odot$  and  $2.0 M_\odot$ , which also provided us with the stellar radii, assuming a simplified mass-radius relation ( $R/R_\odot \simeq (M/M_\odot)$ ). We then drew the parameters of the planetary orbit: the period  $P$ , the planetary radius  $R_p$  (in units of stellar radius), and the impact parameter  $b$ , which determined the orbital inclination. Table 2 details the



**Figure 5.** Binned scatterplot of the obtained Dice coefficient as a function of the S/N for the 90,000 light curves containing transits. The greyscale coding of the scatterplot is visualized by the vertical bar on the right. The marginal distributions of the Dice coefficient and the S/N are also shown as normalized complementary cumulative histograms.

various distributions we used in order to draw all those parameters. The phase was drawn from a uniform distribution.

In total, 100,000 light curves, each containing 20,610 samples, were simulated. For each of the generated light curves, we have injected a transit signal as described above. Thus, we eventually had a total of 200,000 light curves (consisting of pairs of which one contained a transit signal and the other did not). We have split the 100,000 pairs of light curves to 5000 pairs for training, 5000 pairs for validation (used mainly for hyperparameter tuning), and 90,000 pairs for testing. Note that simulating 100,000 light curves was a relatively easy process, and the computational burden was mainly related to the size of the training and validation sets, which is the reason for the larger size of the testing set. The distribution of the signal-to-noise ratio (S/N; as defined in Zucker & Giryes 2018) of the 90,000 time series can be shown in Figure 5, as the complementary cumulative distribution.

## 5. Training

We have trained the U-Net generator and the discriminator simultaneously using the 5000 light curves that contained transits and their binary ground-truth segmentation sequences together, while the classifier used both sets of light curves with and without added transits for 10,000 light curves in total. The global loss function contained coefficients that controlled the relative importance of segmentation (the Dice loss of the generator U-Net) and the evaluation of the discriminator

**Table 3**  
GAN Hyperparameters

Parameter	value
learning-rate	0.0001
beta_1dropout	0.54
beta_2	0.9
batches	10,000
batch-size	32

(adversarial loss). The Dice coefficient was maximized during training, while the adversarial loss of the GAN discriminator was minimized using the Adam optimizer (Kingma & Ba 2014) with hyperparameters shown in Table 3 for 10,000 randomly chosen batches of size 32 out of the 5000 light curves mentioned above.

For the classification training, we used 10,000 light curves with and without added transits. The network was trained for 10,000 batches of 32 data inputs each. This is approximately equivalent to 32 epochs (i.e., going over the whole data set 32 times). Note that each batch is randomly selected from the whole data set, using random permutation of the indices, thus guaranteeing going over all the data. After each batch of training with the generator-discriminator pair, a random batch of 32 light curves (out of the 10,000 mentioned above) was used to train the classifier network using the Adam optimizer, with the hyperparameters shown in Table 4.

**Table 4**  
Classifier Network Hyperparameters

Parameter	value
Learning-rate	0.001
Dropout	0.4
Beta_1	0.9
Beta_2	0.99
Batches	10,000
Batch-size	32

## 6. Results

The ability of a classifier network to identify light curves that contained transits was already demonstrated in Paper I. The main purpose of this work is to present a deep learning approach to identify the transit events in light curves, which had already been labeled to contain transits. Therefore, the results we present here are mainly examples that demonstrate the ability of the neural network to perform this task.

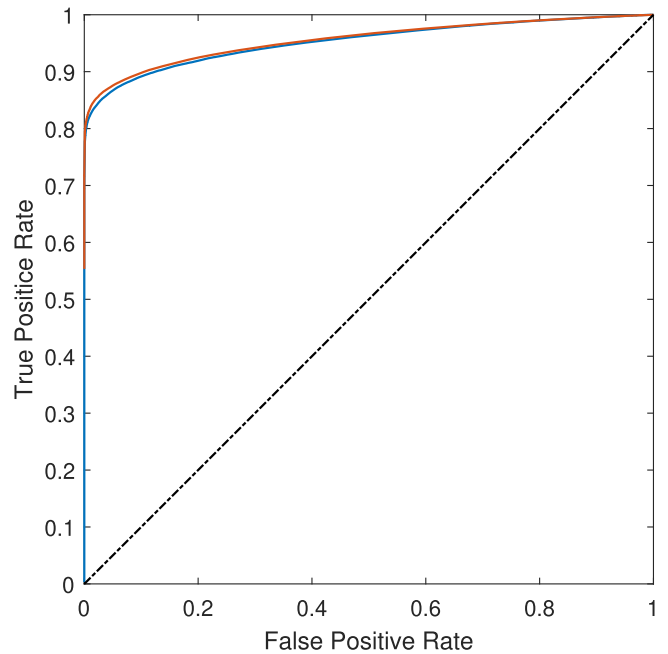
However, we first show in Figure 6 that the classifier does perform satisfactorily in this context, in which it is trained together with the segmentation network. We show this using a receiver operating characteristic (ROC) curve (e.g., Fawcett 2006), which presents the true positive rate as a function of the false-positive rate. The blue curve shows the ROC for the classifier network, when trained alone, separately from the segmentation network. The red curve presents the results after training the two networks together. Clearly the classifier performance is only improving by this combined training, albeit slightly.

Figure 5 shows a binned scatterplot of the dependence of the Dice coefficient on the S/N of the light curves with transits. The marginal distributions of the Dice coefficient and the S/N are shown as normalized histograms representing the complementary cumulative distributions. Clearly, for many light curves the Dice coefficient is close to 1. As can be expected, the performance generally tends to degrade for small S/N, and the performance below a S/N value of 10 becomes quite poor.

Most of the examples we chose to present here had a S/N between 10 and 20 and Dice coefficients higher than 0.9. The last two examples are shown in order to sample S/N values beyond this restricted range. The characteristics of the examples are summarized in Table 5, including the assumed stellar magnitude of the star, the transit main parameters, the S/N and the obtained value of the Dice coefficient.

Example A is a simple and easy to study light curve. As can be seen in the top panel of Figure 7, there is no significant component of red noise—there is almost no apparent long-term nor quasiperiodic variability, and the main source of noise is white noise. This is also evident in the parameters of the GP kernel (Table 5). Note that this star is in the faint end of our simulated sample, which is indeed expected to be dominated by white noise. The transit is deep enough to that the transits can easily be spotted by eye in the light curve. The middle panel shows the injected transit signal at the same scale, while in the lower panel we highlighted by red the samples which the neural network identified as being in transit.

Example B, shown in Figure 8, is of a brighter sixth-magnitude star, and therefore the effect of red noise is supposed to be more significant. The transits are still visible when examining the light curve, but long-term red noise is also easily



**Figure 6.** The performance of the classifier network presented as a ROC curve. Blue: after training the classifier separately from the segmentation network; red: after training the classifier together with the segmentation network. The dotted-dashed line is the so-called “no-discrimination” line, corresponding to a random guess.

seen. Overall, the level of the noise is still low, which allows the very shallow transits to be discernible by eye.

Figure 9 shows Example C, where the red noise is more complicated. The host star is even brighter than the previous example and one can see the many red-noise features in the light curve that are of similar amplitude as the transits. They are mainly caused by the relatively strong quasiperiodic term in the noise with a period of 68 hr, which is close to the transit period of 3.61 days (Table 5). It is not easy to locate the transits by simple visual examination of the light curve.

The quasiperiodic component of the noise is even more prominent in Example D shown in Figure 10. The quasiperiodic noise creates many troughs in the light curve that can be easily mistaken by the human eye to be transits. The neural network still identifies correctly the timing of the periodic transits.

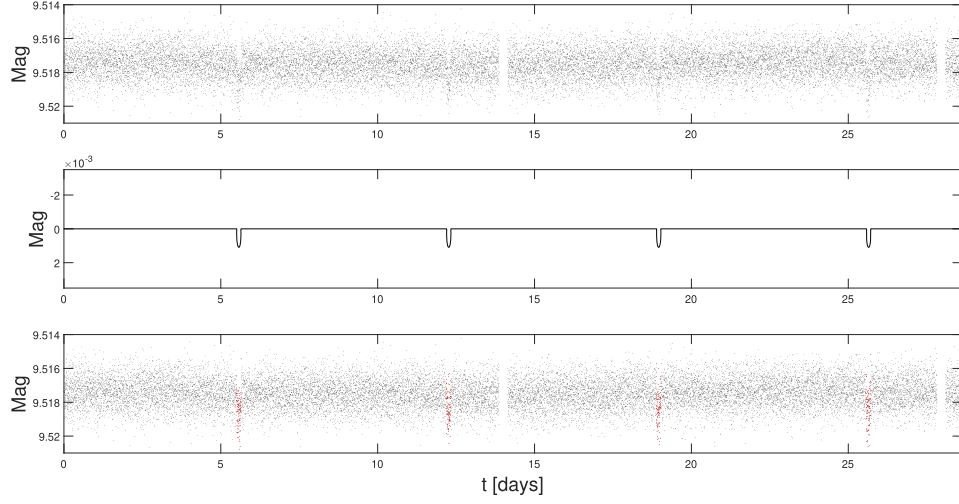
In Figure 11 we show Example E, which demonstrates the ability of the network to deal with the sampling gaps. In this example, both downlink gaps occur during transits of the exoplanet. The coincidence with the gaps does not hamper the ability of the network to perform a correct segmentation and identify the transit even in those times.

Obviously, real life is never perfect, and errors are bound to occur especially for low S/N. Figure 12 shows Example F, which is one of those rare events, where the network identified a transit when none existed. This light curve is dominated mainly by white noise, but with a large amplitude (the star is faint) making for a somewhat lower S/N. The network in this case wrongly identifies a very short transit at some point in day seven. This wrong identification of a transit event was not enough to lower considerably the Dice coefficient.

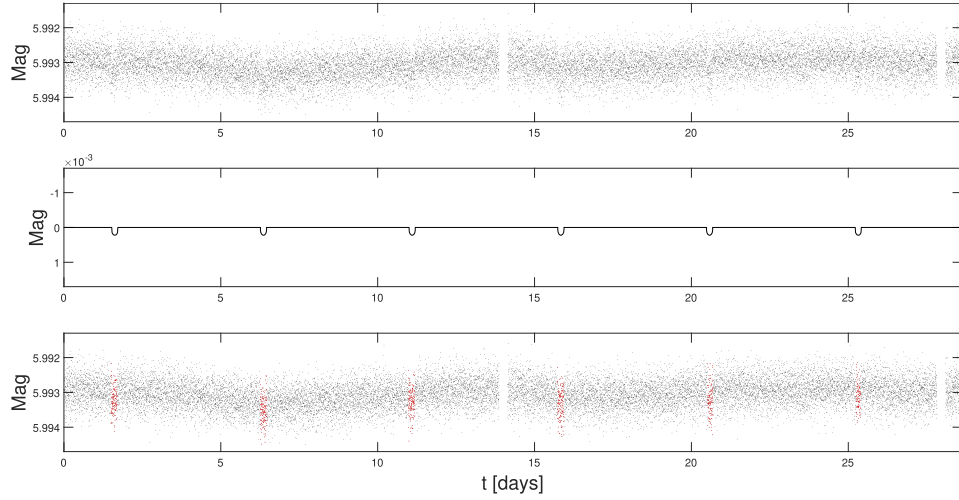
The last two examples demonstrate cases of S/N values beyond the range between 10 and 20, and serve to substantiate the validity of the Dice coefficient as a diagnostic for the segmentation performance. Example G, shown in Figure 13,

**Table 5**  
Details of the Presented Examples

	(mag)	Transit Parameters			GP Parameters						S/N	Dice
		Period (days)	Duration (hr)	Depth (mmag)	$A_w$ ( $\mu$ mag)	$A_s$ ( $\mu$ mag)	$A_q$ ( $\mu$ mag)	$\lambda_s$ (minute)	$T_q$ (hr)	$\lambda_q$ (hr)		
A	9.5	6.69	3.1	1.10	856	26	81	13	103	85	19.4	0.908
B	6.0	4.74	4.4	0.22	361	95	110	3	293	65	16.4	0.901
C	5.4	3.61	3.6	0.24	347	55	104	285	68	71	19.2	0.904
D	5.1	6.26	7.3	0.25	342	7	115	335	11	25	17.5	0.931
E	7.4	4.67	6.0	0.30	441	32	59	7	17	212	18.5	0.917
F	9.0	7.28	4.8	0.61	707	80	99	13	490	58	14.8	0.907
G	6.6	4.49	2.9	0.23	382	42	105	13	361	157	9.5	0.873
H	5.2	6.54	0.38	0.59	344	117	116	36	27	30	28.8	0.029



**Figure 7.** Example A. The upper panel shows the input light curve. The middle panel shows the injected transit, while in the lower panel the light curve is shown again, but the samples that the network tagged as in-transit samples are marked in red. The y-axis scale is identical in all panels to facilitate visual comparison. Four transits can clearly be discerned.



**Figure 8.** Example B. Panels follow the structure of Figure 7. Note the more prominent long-term variability in this specific case.

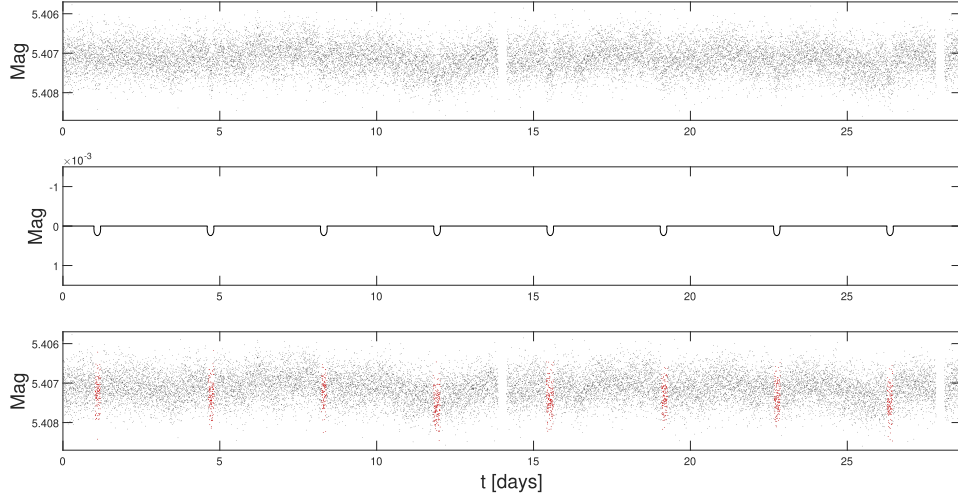
shows a case of low S/N, below 10, which nevertheless exhibits a relatively high Dice coefficient, almost 0.9. This Dice value seems to be perfectly justified based on the satisfactory segmentation. On the other hand, Example H, in Figure 14, shows a case with a high S/N, above 20, but very low Dice coefficient, in line with the very poor segmentation performance as seen in the figure. This poor performance is

probably related to the very short transit duration, combined with a relatively strong presence of red and quasiperiodic noise.

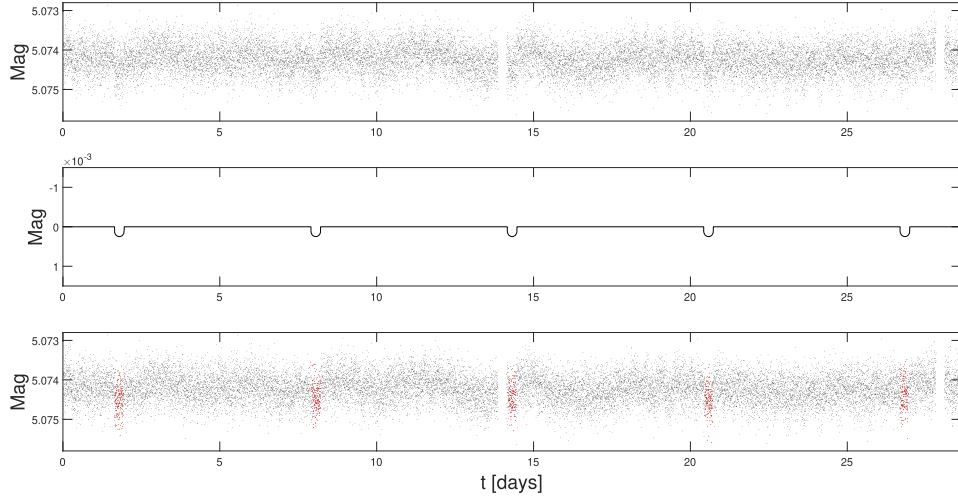
## 7. Discussion

The approach we have described in this paper is meant to complement the detection neural network we have introduced in Paper I, where we have applied a neural network classifier to

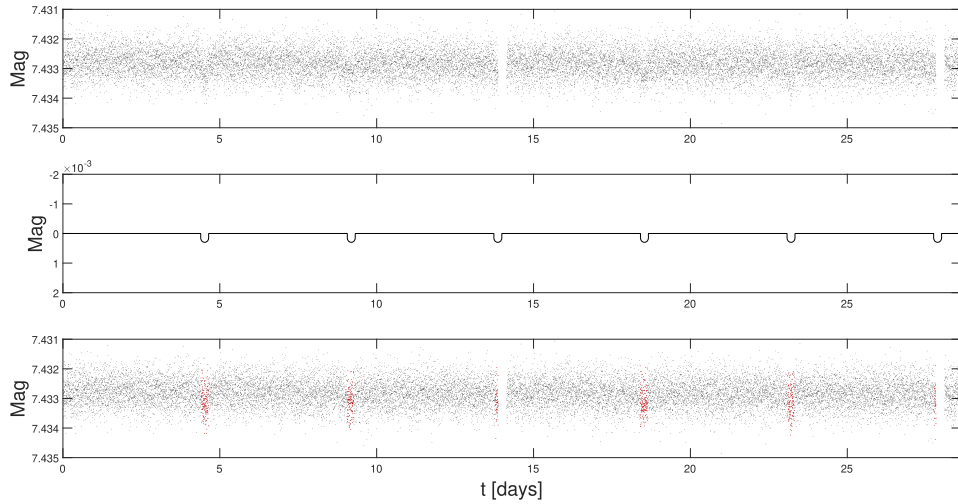




**Figure 9.** Example C. Panels follow the structure of Figure 7. Note the complex patterns of red-noise variability.



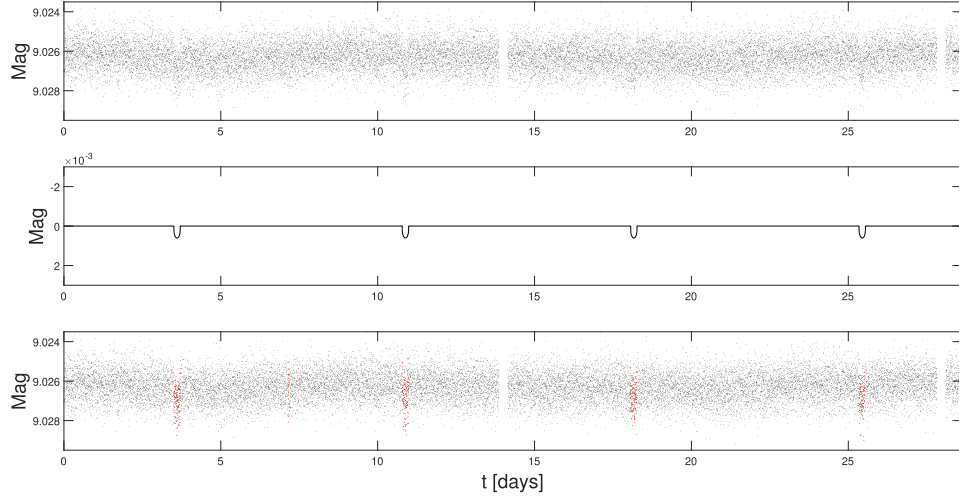
**Figure 10.** Example D. Panels follow the structure of Figure 7. Note the presence of many quasiperiodic transit-like features in the red noise.



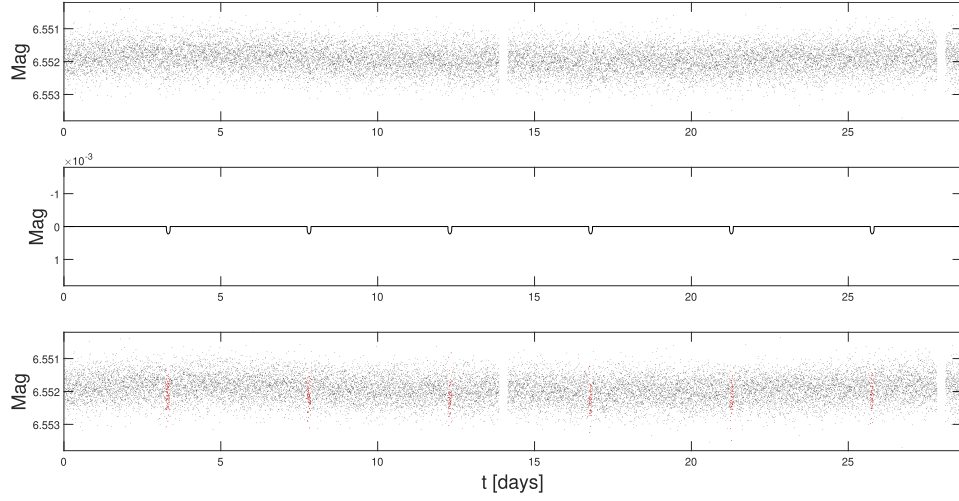
**Figure 11.** Example E. Panels follow the structure of Figure 7. Note the partial coalescence of the third and sixth transits with the downlink gap.

label simulated light curves that contained exoplanetary transits. Before any future analysis of the transits can be attempted, it is necessary to know when the transits actually take place. Only then can they be analyzed for their detailed

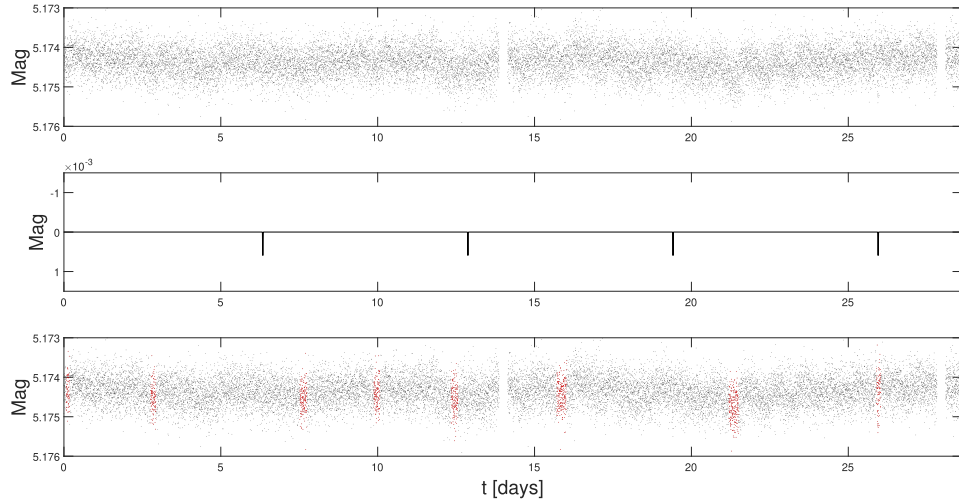
shape, their precise timing, and, most importantly, only then can one vet them and try to infer their nature—whether they are genuinely exoplanetary transits or not. Following the terminology of computer vision, we have dubbed the task of labeling



**Figure 12.** Example F. Panels follow the structure of Figure 7. Note the wrong identification of a short transit around day seven.



**Figure 13.** Example G. Panels follow the structure of Figure 7. All transits were identified in spite of the low S/N, resulting in a high Dice coefficient value.



**Figure 14.** Example H. Panels follow the structure of Figure 7. The segmentation is completely wrong, except for the transit around day 26, as reflected by the very low Dice coefficient.

the in-transit samples “segmentation.” Both tasks, detection and segmentation, become extremely difficult when it comes to very shallow transits in the presence of red noise. We suggest

that deep learning can significantly assist in performing those tasks, as we have tried to demonstrate with the simulations we presented.

The fact that the transits were periodic was very instrumental in using the GAN and the adversarial loss. Among other features of the transit signals, the network apparently learned to favor periodic signals. Once multiple transiting planets exist in the system, the challenge indeed becomes more complicated, but multiple aspects of periodicity still exist in the signal. If additional planets do not transit but only induce TTVs, we estimate that the effect on the segmentation network performance will not be that dramatic, but it still has to be verified. In our future work we intend to tackle those challenges (deviations from pure periodicity) as well.

For simplicity, we have employed in the training of the segmentation network a simple Dice loss that considers equally the in-transit and out-of-transit segments of the light curve. Assigning different weights to those areas may improve the performance and avoid the rare events of erroneous labels, like the one we show in our last example (Figure 12). Furthermore, we note that the network usually reproduces quite well the phases of the transits, but the exact timing of the ingress and egress tend to suffer some more inaccuracies, probably caused by interfering features of the red noise. This might also be corrected by a more sophisticated loss function that would take into account the shape of the signal and the timing of the ingress/egress stages.

Unlike the network we have introduced in Paper I, the output of the network is not a binary decision regarding the presence of a transit signal, but a sequence of decisions tagging the original samples as in- or out-of-transit. The network itself outputs a sequence of real numbers, and a threshold is applied to those numbers to produce the binary segmentation sequence. In principle, a similar mechanism can be applied to perform detrending, in a way that would not destroy the transit signal. This will be another central aim of our future efforts—detrend the signals and remove the red noise, without affecting the transit signal. This would be an extremely challenging task.

We have trained the neural network such that it would tag transits that are really on the border of detection. The PLATO mission will be targeting exactly such events. Deep learning techniques like the one we have presented here are bound to play a significant role in those efforts. Transits such as those presented in Figures 7 and 8 might be vetted reasonably well using current techniques. However, transits such as those shown in our next examples, the more difficult ones, might require additional observations, photometric, spectroscopic, or others that still have to be thought of.

In addition, similar to Paper I, we added a binary classifier that uses information from the residual connections of the generator (Figure 4) to determine whether the light curve seems to contain a transiting planet signal, so that our final network performs the full task of identifying light curves with periodic transits and performing the segmentation.

Deep learning is an exploding field of data science. The exoplanetary community already acknowledges that and tries to use those techniques in its endeavor. However, most of the efforts focus around vetting transits using various variants of CNNs. The work we have presented here may serve as a reminder that there are other flavors of neural networks and training methods like U-Nets and GANs, which may also be very useful to exoplanet research.

This work was supported by a grant from the Tel Aviv University Center for AI and Data Science (TAD) and by the

Ministry of Science, Technology and Space, Israel. R.G. acknowledges support by ERC-stg SPADE (grant No. 757497).

## ORCID iDs

Shay Zucker  <https://orcid.org/0000-0003-3173-3138>  
Raja Giryes  <https://orcid.org/0000-0002-2830-0297>

## References

- Aigrain, S., Parviainen, G., & Pope, B. J. S. 2016, *MNRAS*, **459**, 2408  
 Ansdell, M., Ioannou, Y., Osborn, H. P., et al. 2018, *ApJL*, **869**, L7  
 Arjovsky, M., Chintala, S., & Bottou, L. 2017, in *ICML'17: Proc. of the 34th Int. Conf. on Machine Learning*, ed. D. Precup & Y. W. Teh (Brookline, MA: Microtome), 214, <https://proceedings.mlr.press/v70/arjovsky17a.html>  
 Blau, Y., & Michaeli, T. 2018, in 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), ed. D. Forsyth et al. (Los Alamitos, CA: IEEE Computer Society), 6228  
 Borucki, W. J., Koch, D., Basri, G., et al. 2010, *Sci*, **327**, 977  
 Dattilo, A., Vanderburg, A., Shallue, C. J., et al. 2019, *AJ*, **157**, 169  
 Deleuil, M., Moutou, C. & the CoRoT Exoplanet Science Team 2010, in *Physics and Astrophysics of Planetary Systems (EAS Publication Series Vol. 41)* ed. T. Montmerle, D. Ehrenreich, & A.-M. Lagrange (Les Ulis: EDP Sciences), 85  
 Dice, L. R. 1945, *Ecology*, **26**, 297  
 Dvash, E., Peleg, Y., Zucker, S., & Giryes, R. 2022, *StrudelTAU/ShallowTransitsDL*, v1.0.0, Zenodo, doi:[10.5281/zenodo.6304556](https://doi.org/10.5281/zenodo.6304556)  
 Fawcett, T. 2006, *PaReL*, **27**, 861  
 Goodfellow, I. J., Bengio, Y., & Courville, A. 2016, *Deep Learning* (Cambridge, MA: MIT Press)  
 Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., et al. 2014, in *Proc. of the 27th Int. Conf. on Neural Information Processing Systems (NIPS 14)*, 2, ed. Z. Gharhamani et al. (Cambridge, MA: MIT Press), 2672, <https://dl.acm.org/doi/10.5555/2969033.2969125>  
 Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., & Courville, A. C. 2017, in *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, ed. I. Guyon et al. (Red Hook, NY: Curran Associates, Inc), 5767, <https://proceedings.neurips.cc/paper/2017/file/892c3b1c6dcd52936e27cbd0ff683d6-Paper.pdf>  
 He, K., Zhang, X., Ren, S., & Sun, J. 2016, in 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), ed. L. Agapito et al. (Los Alamitos, CA: IEEE Computer Society), 770  
 Howarth, I. D. 2011, *MNRAS*, **418**, 1165  
 Jenkins, J. M., Tenenbaum, P., Caldwell, D. A., et al. 2018, *RNAAS*, **2**, 47  
 Kingma, D. P., & Ba, J. 2014, arXiv:[1412.6980](https://arxiv.org/abs/1412.6980)  
 Kovács, G., Zucker, S., & Mazeh, T. 2002, *A&A*, **391**, 369  
 Kreidberg, L. 2015, *PASP*, **127**, 957  
 LeCun, Y., Bengio, Y., & Hinton, G. 2015, *Natur*, **521**, 436  
 LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. 1998, *IEEEP*, **86**, 2278  
 Long, J., Shelhamer, E., & Darrell, T. 2015, in 2015 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), ed. K. Grauman et al. (Los Alamitos, CA: IEEE Computer Society), 640  
 Milletary, F., Navab, N., & Ahmadi, S.-A. 2016, in *Proc. of the 4th Int. Conf. on 3D Vision (3DV)*, IEEE 2016, ed. N. Snavely et al. (Los Alamitos, CA: IEEE Computer Society), 565  
 Nair, V., & Hinton, G. E. 2010, in *Proc. of the 27th Int. Conf. on Machine Learning (ICML-10)*, ed. J. Fürnkranz & T. Joachims (Madison, WI: Omnipress), 807  
 Osborn, H. P., Ansdell, M., Ioannou, Y., et al. 2019, *A&A*, **633**, A53  
 Rauer, H., Aerts, C., Cabrera, J., et al. 2016, *AN*, **337**, 961  
 Ricker, G. R., Winn, J. N., Vanderspeck, R., et al. 2015, *JATIS*, **1**, 014003  
 Ronneberger, O., Philipp, F., & Thomas, B. 2015, in *Medical Image Computing and Computer-Assisted Intervention (MICCAI 2015)*, ed. N. Navab et al. (Berlin: Springer), 234  
 Rumelhart, D. E., Hinton, G. E., & Williams, R. J. 1986, *Natur*, **323**, 533  
 Schmidhuber, J. 2015, *NN*, **61**, 85  
 Shallue, C. J., & Vanderburg, A. 2018, *AJ*, **155**, 94  
 Sørensen, T. 1948, *Biol. Skr. Dan. Vid. Sel.*, **5**, 1  
 Yu, L., Vanderburg, A., Huang, C., et al. 2019, *AJ*, **158**, 25  
 Zucker, S., & Giryes, R. 2018, *AJ*, **155**, 147