

# Nigraha: Machine-learning-based pipeline to identify and evaluate planet candidates from TESS

Sriram Rao <sup>1</sup>★, Ashish Mahabal <sup>2,3</sup>, Niyanth Rao<sup>4</sup> and Cauligi Raghavendra<sup>5</sup>

<sup>1</sup>LinkedIn Corporation, Mountain View, CA 94043, USA

<sup>2</sup>Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA

<sup>3</sup>Center for Data Driven Discovery, California Institute of Technology, Pasadena, CA 91125, USA

<sup>4</sup>Saratoga High School, Saratoga, CA 95070, USA

<sup>5</sup>University of Southern California, Los Angeles, CA 90089, USA

Accepted 2021 January 13. Received 2020 December 28; in original form 2020 October 9

## ABSTRACT

The Transiting Exoplanet Survey Satellite (TESS) has now been operational for a little over two years, covering the Northern and the Southern hemispheres once. The TESS team processes the downlinked data using the Science Processing Operations Center (SPOC) pipeline and Quick Look pipeline (QLP) to generate alerts for follow-up. Combined with other efforts from the community, over 2000 planet candidates have been found of which tens have been confirmed as planets. We present our pipeline, *Nigraha*, that is complementary to these approaches. *Nigraha* uses a combination of transit finding, supervised machine learning, and detailed vetting to identify with high confidence a few planet candidates that were missed by prior searches. In particular, we identify high signal-to-noise ratio shallow transits that may represent more Earth-like planets. In the spirit of open data exploration, we provide details of our pipeline, release our supervised machine learning model and code as open source, and make public the 38 candidates we have found in seven sectors. The model can easily be run on other sectors as is. As part of future work, we outline ways to increase the yield by strengthening some of the steps where we have been conservative and discarded objects for lack of a datum or two.

**Key words:** methods: data analysis – techniques: photometric – planets and satellites: detection – planetary systems.

## 1 INTRODUCTION

Launched in 2018, the Transiting Exoplanet Survey Satellite (*TESS*) mission (Ricker et al. 2014) is an all-sky digital survey aimed at observing planetary transits of their host stars. The *TESS* spacecraft monitors  $24^\circ \times 96^\circ$  sectors of the sky for contiguous periods of 27 d apiece; data is downlinked once every 13.5 d for further processing. *TESS* monitors the sky in an anti-Sun direction, allowing prompt ground-based follow-up if the candidates are identified quickly. For context, the data for each sector comprises approximately 20 000 light curves. On average, less than 1 per cent of these are potential planetary candidates for follow-up. Effectively, the number of planet candidates are dwarfed by a vast majority of false positives (i.e. variable stars, eclipsing binaries, etc.) and/or are affected by instrument systematics which cannot be filtered out.

Classical machine-learning-based methods such as decision trees,  $k$ -nearest neighbors, and random forests have been explored for classifying light curves (McCauliff et al. 2015; Thompson et al. 2015; Armstrong, Pollacco & Santerne 2016; Armstrong et al. 2018). Of these systems, the *Autovetter* (McCauliff et al. 2015), which is a random forest classifier that makes decisions based on the metrics generated by the Kepler pipeline, has been successfully used to

make initial dispositions for Kepler/K2 candidates. Convolutional neural networks (CNNs) that are effective at image recognition tasks have been explored as an alternate approach for classifying light curves (Pearson, Palafox & Griffith 2017; Ansdell et al. 2018; Shallue & Vanderburg 2018; Zucker & Giryes 2018; Dattilo et al. 2019; Schanche et al. 2019; Yu et al. 2019; Osborn et al. 2020). These models use supervised learning with a CNN where the neural network learns the characteristics from light curves. At a high level, these systems employ the following procedure for rapid identification of targets: (i) pre-process the light curves to search for transits (i.e. periodic dips in the flux) to derive transit parameters such as period and depth, (ii) from the transit parameters build phase-folded light curves which are subsequently used to train a CNN-based model with labeled data, and (iii) make predictions on newly observed data. This procedure has been successfully used in practice with Kepler and K2: (i) Shallue & Vanderburg (2018) use their *Astronet* deep learning model on Kepler data to identify several candidates from which they validated two new exoplanets and (ii) Dattilo et al. (2019) have adapted *Astronet* model for K2 data (dubbed, *Astronet-K2*) to discover two new exoplanets. Ansdell et al. (2018) extend the basic *Astronet* model to include stellar parameters and demonstrate with Kepler data that their *Exonet* model improves model accuracy by using domain information.

In this paper, we explore the use of CNNs for classifying *TESS* light curves to search for planet candidates. A natural starting

\* E-mail: sriramrao@gmail.com (SR); aam@astro.caltech.edu (AM)

point is these existing models based on Kepler data. However, the systematics with individual TESS sectors are significantly different when compared to Kepler/K2 (*viz.*, 27 d of data versus 4 yr/81 d of data). Hence, prior deep-learning approaches for classifying TESS data such as Yu et al. (2019), Osborn et al. (2020) adapt and modify *Astronet* and *Exonet*, respectively, and identify new candidates. Our work is similar in spirit in that we start with a hybrid of these two models and develop novel extensions that allow us to explore different aspects of the input parameter space to identify candidates.

The core contribution of this paper is in the design and evaluation of a deep learning-based pipeline for identifying planetary candidates from TESS data. In building our pipeline, *Nigraha*,<sup>1</sup> two aspects are novel, namely the input representation of the light curve to the CNN model (Section 2) and the model’s architecture. We train and evaluate the model using labeled data from actual TESS sectors (Section 3). Our model has an AUC (area under the curve) score of 0.908. We also evaluate our model by augmenting our test set with bona fide planets and find that it recovers 91.9 per cent (91 of 99) (Fig. 7). We apply our trained model to make predictions on data from TESS Sector 6 and Sectors 21–26 and find new candidates (Section 4). We then perform preliminary vetting on a subset of these candidates using a series of steps, including *DAVE* (Kostov et al. 2019a), *Gaia* Renormalised Unit Weight Error (RUWE), and *2MASS*, to identify a promising sample of 38 new candidates for follow-up (Table 3, 4 and 5). Finally, Section 5 offers concluding remarks and presents directions for future work.

## 2 DATA PREPARATION

In this section, we describe the data preparation steps for generating input representation of light curves to the machine learning models. We describe our procedure for searching for transits in light curves (Section 2.3), and then use that for building the input representation (Section 2.4). We begin by describing the labeled data sets used in this paper.

### 2.1 Data sets

We use pre-processed TESS light curves downloaded from Mikulski Archive for Space Telescopes (MAST).<sup>2</sup> These have a 2-min cadence over a 27-d period and have instrumental systematics removed. We use the Pre-Search Data Conditioning Simple Photometry flux (PDCSAP\_FLUX) time series from each light curve to search for *threshold crossing events* (TCEs). These are periodic flux variations potentially caused by exoplanets transiting their host star. We train our models using TCEs from various TESS sectors and then use the trained models to identify new candidates.

### 2.2 Training labels

We use two sets of labeled data for this work. We use the TESS TOI catalogue<sup>3</sup> to construct our first labeled data set (**DS1** in Table 1). This catalogue specifies the TESS identifier for which an alert was generated (and the associated light curve is available from MAST), the disposition for that target, and additional metadata. With TESS labels, KP corresponds to planets known prior to launch, CP are

**Table 1.** Training labels for the TCEs from the two data sets used in this paper. In the labels column, P indicates positive labels, and N, the negative labels.

Type	Count		Label	
	DS1	DS2	P	N
Known planet (KP)	166		✓	
Confirmed planet (CP)	38		✓	
Planet candidate (PC)	274	165	✓	
Eclipsing binary (EB)		711		✓
False positive (FP)	139			✓
Instrumental noise (IS)		520		✓
Variable star (V)		656		✓
Other (O)	1	1		✓
Junk (J)		3917		✓
Total	618	5970	643	5945
Grand total		6588		

confirmed planets originally identified using TESS data, and PC are targets that have been identified as planet candidates. Note that the TESS TOI catalogue contains alerts/labels for targets in all TESS sectors to date. The catalogue is therefore a ‘living’ document. When the first draft of this paper was complete, the TOI catalogue had labels from Sectors 1 to 26. From this data set, we set aside labeled targets from sectors on which we make predictions. That is, in this paper we train our models and then generate predictions for planetary candidates in Sector 6/Sectors 21–26. The second labeled data set is from Yu et al. (2019) (**DS2** in Table 1) which they have made publicly available<sup>4</sup> and we use without any modifications. Their labels are for targets observed from TESS Sectors 1–5. For TCEs common in the two data sets, we use the labels from **DS1** as ‘ground truth’. The dispositions for the various TCEs and the data sets from which they were obtained is shown in Table 1. Note that the data set is imbalanced: 9.8 per cent correspond to positive labels.

### 2.3 Transit search

Searching for transits in a stellar light curve entails computing a periodogram from which transit parameters such as, period ( $P$ ), duration ( $t_{\text{dur}}$ ), and epoch ( $T_0$  or mid-point of transit) are obtained. Traditionally, BOX LEAST SQUARES (BLS; Kovács, Zucker & Mazeh 2002) has been the widely used algorithm for this purpose. It builds a periodogram by searching for box-like transits in a given light curve. Recently, TRANSIT LEAST SQUARES (TLS; Hippke & Heller 2019) has been proposed as an alternative to BLS with the objective of improving the detection of shallow transits i.e. sensitivity towards detection of transits caused by near-Earth-sized planets. The key idea in TLS is to build an analytical model of the transit with stellar limb darkening (Mandel & Agol 2002). Hippke & Heller (2019) point out that TLS is able to detect shallow transits at an acceptable computational cost (i.e. approximately 10 s per TESS light curve when run on an eight-core Intel i9 laptop). Furthermore, the false positives produced by TLS are comparable in number with BLS when detecting transits. When applied to Kepler/K2 data, Heller, Rodenbeck & Hippke (2019a), Heller, Hippke & Rodenbeck (2019b) show that TLS is able to discover Earth-sized planets that were previously missed. That is, BLS did not recover the transits of a sub-Earth in a multiplanet system on K2SFF detrended light curves; in contrast, TLS did. These

<sup>1</sup> Sanskrit word to mean *determined* with part of the word meaning a planet.

<sup>2</sup> [http://archive.stsci.edu/tess/bulk\\_downloads/bulk\\_downloads.f-fi-tp-1c-dv.html](http://archive.stsci.edu/tess/bulk_downloads/bulk_downloads.f-fi-tp-1c-dv.html)

<sup>3</sup> <https://exofop.ipac.caltech.edu/tess/>

<sup>4</sup> <https://github.com/yuliang419/Astronet-Triage/blob/master/astromet/tces.csv>

results from TLS are particularly interesting in the context of TESS: the candidate set of planets targeted by TESS include super-Earths orbiting close to their host star for which the transits are necessarily shallow. Motivated by these observations, we use TLS for transit detection.

For all the data sets used in this paper, we compute transit parameters using the publicly available open source version of TLS (TLS version 1.0.25). As a preprocessing step we detrend the PDCSAP\_FLUX time series from each TESS light curve file using an approach similar to Heller et al. (2020). We remove variability on timescales longer than the maximal expected transit duration with Tukey’s biweight filter using the WOTAN software package (Hipke et al. 2019). WOTAN provides a convenience API for estimating the transit duration as a function of mass, radius, and period. To this API, we input mass, radius estimates from the TESS input catalogue (Stassun et al. 2018) and as a prior, we set the maximum period to 13.5 d (i.e. half the duration of the observing period of a TESS sector). We then use a running window of three times this duration for detrending the light curve. For a more detailed discussion on the choices of the filter and the associated parameters, we refer the reader to Section 2.1 of Heller et al. (2020). Subsequently, we use TLS APIs to compute the power spectrum on the detrended data. For this step, we input stellar limb darkening parameters for TESS targets as estimated by Claret (2018) and use default values for the remaining parameters in TLS APIs (see documentation<sup>5</sup>). Based on these inputs, TLS computes the transit parameters corresponding to the best fit for the data. Additionally, TLS also returns the number of transits, signal-to-noise ratio (SNR) as defined by Pont, Zucker & Queloz (2006), and the signal detection efficiency (SDE) which represents the peak in the TLS power spectrum. Using these values, for a light curve to contain a TCE, we require (i) at least two transits, (ii) an SNR > 7.1, and (iii) an  $SDE_{\text{TLS}} > 9$ .

The threshold applied on  $SDE_{\text{TLS}}$  impacts TLS’s false positive rate of recovering transits. It represents a trade-off between the false positive rate and the number of TCEs that are recovered by the transit search. Our choice of  $SDE_{\text{TLS}} > 9$  is based on the TLS false positive rate of recovering transits to about  $10^{-4}$  in the limiting case of white noise [see Section 3.1 of Heller et al. (2019a)]. Of the 1168 light curves we searched from the TOI catalogue, we recover 618 TCEs that are in **DS1**. An interesting observation is that our procedure recovers vast majority of the TCEs with a final disposition of KP, CP (see Fig. 2). Our choice for  $SDE_{\text{TLS}}$  threshold is intentionally conservative since machine vetting for TESS candidates is presently limited (see Section 4.3). With even better vetting, we could, in future, lower the SDE threshold to recover more of the KPs, CPs (and PCs).

Besides computing  $P$ ,  $t_{\text{dur}}$ ,  $T_0$ , TLS also returns several other transit parameters that we use and they are described below (see Section 2.5).

## 2.4 Input representation

We use the LIGHTKURVE toolkit (Lightkurve Collaboration 2018) to prepare the light curves as input to our machine learning model (see Section 3). Our method is similar to Shallue & Vanderburg (2018). We iterate over each light curve containing a TCE and perform the following sequence of steps. First, we clean the light curve to remove outliers that are at least  $5\sigma$  above the median. Secondly, we remove low frequency trends from the light curve on timescales longer than three times the expected transit duration using the Savitzky–Golay

filter. We use LIGHTKURVE’s `flatten()` API. To prevent the in-transit points from being detrended out, from the mid-point of each transit in the complete light curve, we mask out points that are within  $-1.5 \cdot t_{\text{dur}}$  to  $1.5 \cdot t_{\text{dur}}$ . We use default values for the remaining parameters<sup>6</sup> of this API (see documentation<sup>7</sup>). Thirdly, we phase-fold the detrended light curve at the epoch (i.e.  $T_0$  that we computed above) so that the mid-point of the transit in all the light curves is aligned. Finally, we bin the phase-folded light curves to construct two views: (i) a ‘global view’ with 201 bins which corresponds to a transit event for a single period (i.e. the phase between -0.5 and 0.5), and (ii) a ‘local view’ with 81 bins which is a ‘zoom’ in view of the transit between  $-2.0 \cdot t_{\text{dur}}$  to  $2.0 \cdot t_{\text{dur}}$ . In Fig. 1, the columns labeled ‘Global view’ and ‘Local view’ show sample representation of the respective global/local views for the various classes of light curves.

With just the global/local views as input, a neural network can be trained to classify whether a light curve contains a transit. However, further separation of eclipsing binaries from genuine planet candidates is non-trivial. This is partly because the neural network’s input data representation is focused on the characteristics of a single transit event and the shape of that transit event in the light curves of those two systems is quite similar. Hence, to improve model performance, we incorporate, with certain modifications, additional secondary views as input as suggested in Shallue & Vanderburg (2018).

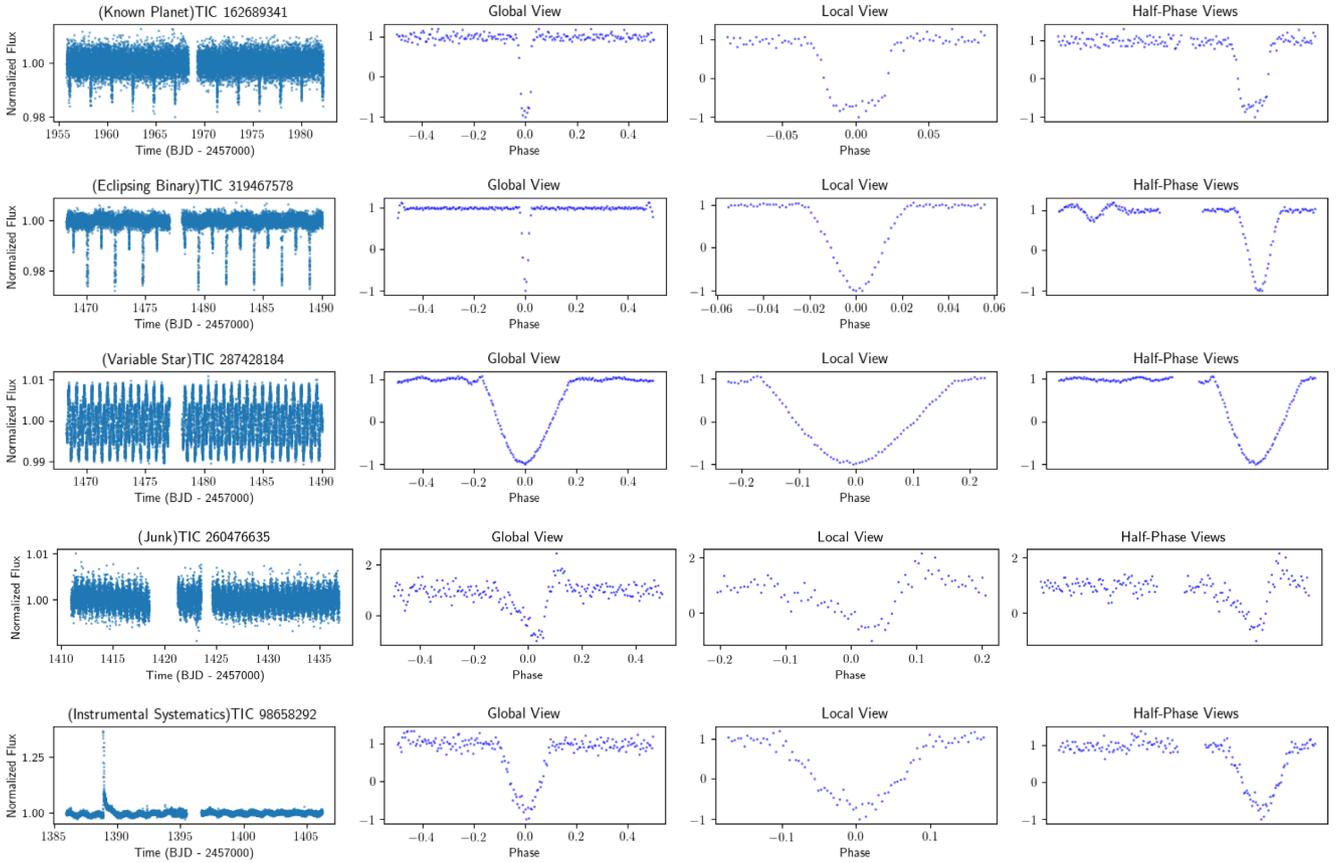
Our approach for constructing a secondary view is based on the characteristics of light curves of eclipsing binaries: The light curve contains a secondary eclipse whose location is dependent on the system’s orbit. Specifically, for eclipsing binaries in a circular orbit, the secondary eclipse occurs at half the period (i.e. midway between a pair of consecutive primary transits). Based on this observation, we build an auxiliary input representation that enables our model to prune a class of eclipsing binaries from planet candidates. We construct a secondary view (called ‘Half-Phase’ view) that is similar to the local view (described above), but with a few changes: (i) we bin and phase-fold the light curve at  $T_0$  (i.e. phase = 0), and (ii) we bin and phase-fold the light curve at  $T_0 + 0.5 \cdot P$  (i.e. phase = 0.5), and (iii) we concatenate these two and normalize them as before. We illustrate these ideas in Fig. 1. For an eclipsing binary (second row in Fig. 1), our representation includes the deep secondary eclipse at phase 0.5. Hence, the Half-Phase view comprises the primary and a secondary eclipse. In contrast, for a bona fide planet (first row in Fig. 1), there is no secondary eclipse, and hence the Half-Phase view contains only the primary transit.

Our method differs from previously explored methods for constructing secondary views. Yu et al. (2019) construct a secondary view by searching the light curves for the mid-point of the most likely secondary eclipse (i.e. masking the transits from the light curve and then using BLS to fit a box to find a transit) and then generate a binned/phase-folded view as input to their `ASTRONET`-vetting model. While this approach will likely identify secondary eclipses occurring anywhere within a period, a downside is that it may not generalize well. For instance, planetary systems with multiple transiting planets could be incorrectly classified as eclipsing binaries, adversely impacting model performance (i.e. low recall). Hence, we chose to pursue a different approach

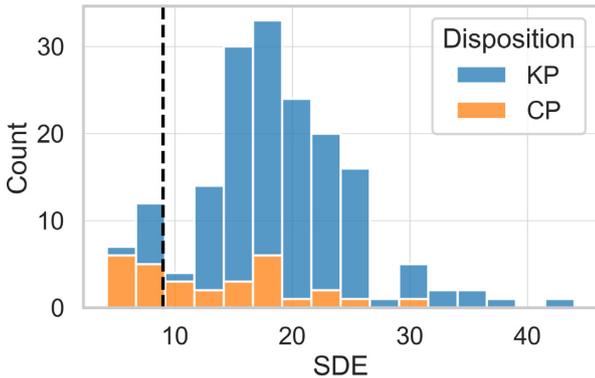
<sup>6</sup>`window_length = 101, polyorder = 2,`  
`break_tolerance = 5, niters = 3, sigma = 3`

<sup>7</sup><https://docs.lightkurve.org/>

<sup>5</sup>[https://transitleastssquares.readthedocs.io/\\_/downloads/en/latest/pdf/](https://transitleastssquares.readthedocs.io/_/downloads/en/latest/pdf/)



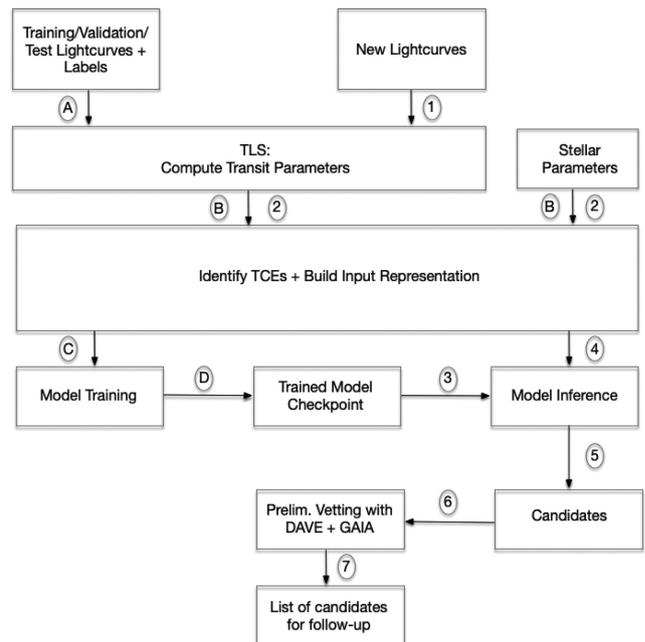
**Figure 1.** Sample views for the various classes of TCEs that are input to our binary model. The classes, in order from top to bottom, are KP, EB, V, J, and IS. The y-axes for the three views are normalized fluxes. Only KPs are used as positive labels in our model, and the rest as negative labels.



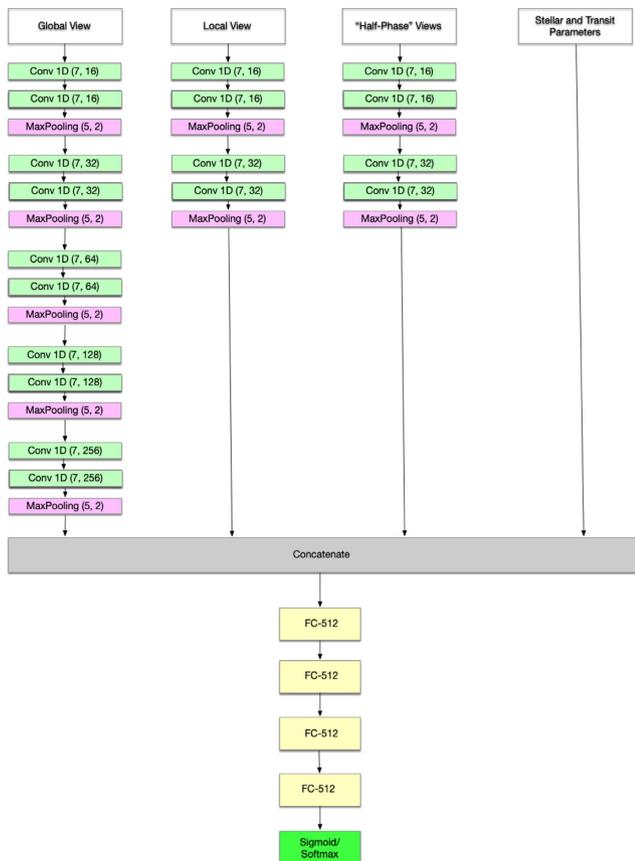
**Figure 2.** Distribution of the TLS computed SDEs for the various KP, CP from the TOI catalogue. Our choice of a threshold for a TCE ( $SDE_{\text{TLS}} > 9$ , corresponding to the vertical dotted line) is conservative and hence, we fail to recover 29 of these TCEs.

## 2.5 Stellar/transit parameters

Ansdeell et al. (2018) show that providing additional domain features can improve model performance for deep learning-based classifiers. Therefore, similar to Osborn et al. (2020), we provide the following additional inputs to our model: (i) stellar parameters that we download from MAST catalogues namely stellar effective temperature ( $T_{\text{eff}}$ ), stellar log  $g$  ( $\log g$ ), star radius ( $R_*$ ), star mass, luminosity, stellar density, and (ii) transit parameters computed by TLS namely,



**Figure 3.** Nigraha pipeline for model training, inference, and preliminary vetting of TCEs.



**Figure 4.** Neural network architecture for our model. Convolutional layers are denoted as Conv1D  $\langle$ kernel\_size, number of filter maps $\rangle$ , pooling layers are denoted as MaxPooling  $\langle$ pool\_size, stride\_size $\rangle$ , and fully connected layers are denoted as FC- $\langle$ number of units $\rangle$ . The activation function in the output layer depends on the type classifier – sigmoid for binary and softmax for multiclass classification. Currently we use only a binary model.

transit depth, transit duration,  $R_p/R_s$ , mean odd/even transit depth. Lastly, similar to Osborn et al. (2020), we normalize each of these additional data columns by subtracting the median value and dividing by the standard deviation.

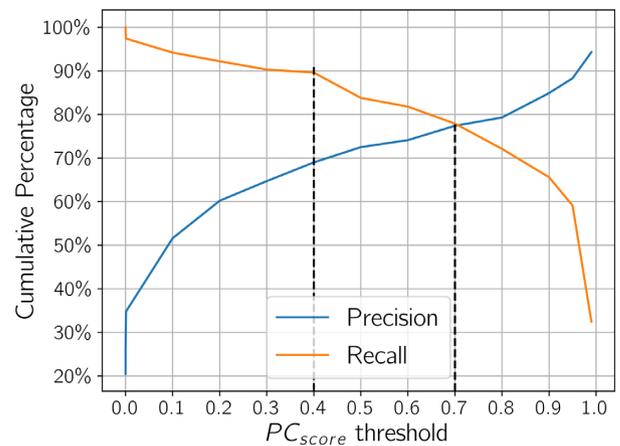
### 3 NIGRAHA PIPELINE

We begin this section by providing an overview of our *Nigraha* pipeline for identifying planetary candidates (Section 3.1). Next, we describe the neural network architecture of our machine learning model (Section 3.2). We then describe our model training and prediction methodology (Section 3.3). We then evaluate model performance in (Section 3.4) and provide a brief comparison with prior approaches (Section 3.5).

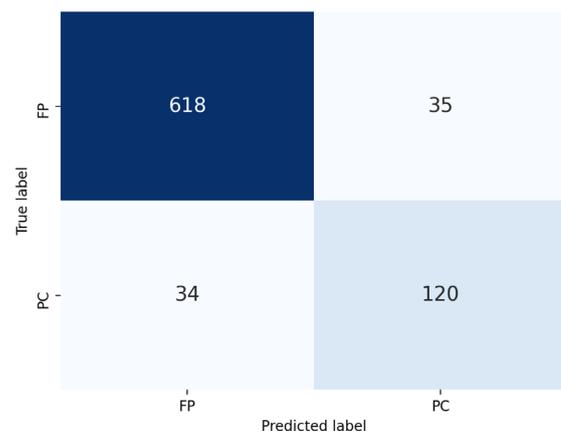
#### 3.1 Nigraha pipeline overview

Fig. 3 depicts the training and inference (i.e. prediction) flows in our pipeline. Steps (A)–(D) comprise the *model training* path.

- (i) The input light curves are processed using TLS which computes transit parameters.
- (ii) The transit parameters are used to pre-process the TCEs to build the Global/Local/Half-Phase views. Also, the stellar parameters are combined to generate the input representation for the TCEs.



**Figure 5.** Precision-recall curves as a function of  $PC_{score}$ . Moving the threshold leads to a precision-recall trade-off.



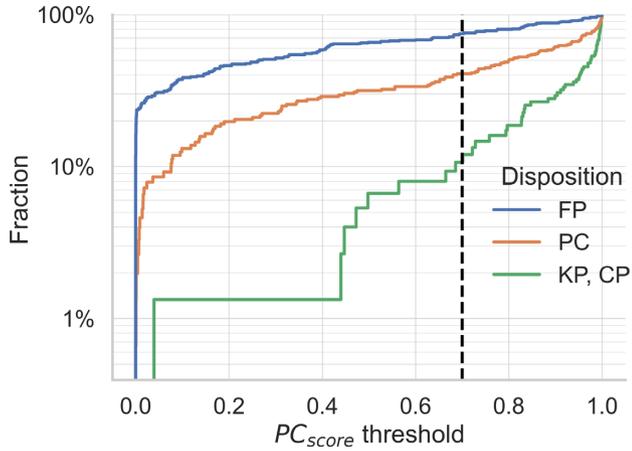
**Figure 6.** Confusion matrix for our model where  $PC_{score} > 0.7$  is used to identify planet candidates. PC are planet candidates (i.e. positive instances), and FP are false positives (i.e. negative instances).

- (iii) The input representations are then used for model training.
- (iv) At the end of model training, a checkpoint is generated which stores the model weights.

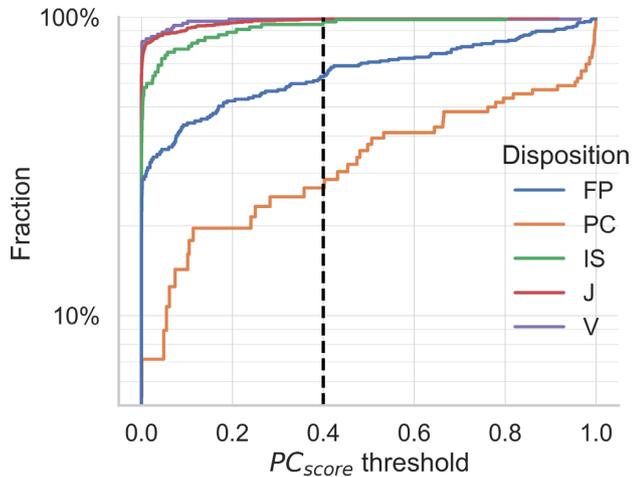
Steps (1)–(4) comprise the *prediction* path. Similar to the training path, Steps (1) and (2) generate the input representation for the TCEs. Next, the model checkpoint is loaded (Step 3) and predictions are generated for each input (Step 4). The inference step generates a prediction (i.e. a score  $\in (0, 1)$ ) for each TCE and this score ( $PC_{score}$ ) can be used to rank the candidates. We apply a threshold on  $PC_{score}$  to choose candidates (Step 5). We perform preliminary vetting of the candidates (Step 6). This is the final step of the pipeline and the output produced by this step is a list of candidates for follow-up (Step 7).

#### 3.2 Model architecture

The *Nigraha* pipeline is built around a convolutional neural network based deep learning model. The architecture of our model is shown in Fig. 4. We use the *Astronet* architecture (Shallue & Vanderburg 2018) as the starting point for our model and the convolutional layers for global/local views is similar. We then extend the model to include the half-phase view which has layers identical to the local view. The Global, Local, and Half-Phase views are input as separate



**Figure 7.** A plot of the CDF that shows the proportion of TCEs scored by our model as a function of the  $PC_{score}$  threshold. In particular, as illustrated by the vertical dotted line, Nigraha assigns a high score ( $PC_{score} > 0.7$ ) to nearly 90 per cent of the bona fide planets (KP, CP).

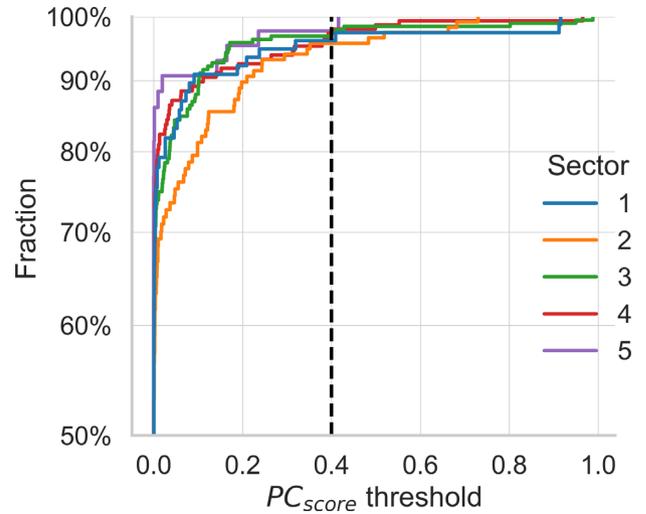


**Figure 8.** A CDF plot that shows how the false positive TCEs are scored by our model as a function of the  $PC_{score}$  threshold. For comparison, we also include the CDF of the scores for the PC class.

convolutional columns. Similar to Exonet model (Ansdell et al. 2018), the outputs of these convolutional layers are concatenated with the stellar and transit parameters and then passed to the fully connected layers. We use the `relu` activation function for the convolutional layers. Our model can be configured to perform either binary or multiclass classification. For binary classification, the output layer uses a `sigmoid` activation function to generate a score in range (0, 1) that corresponds to the likelihood of the TCE being a planet candidate. That is, values closer to 1 indicate a high likelihood of being a PC whereas values closer to 0 indicate that the TCE is likely to be a FP. Extending our model to support multiclass classification is part of on-going work in this project and is beyond the scope of this paper. We defer discussion to Section 5.

### 3.3 Model training and prediction methodology

We implement our deep learning model using Keras (Chollet et al. 2015) on Tensorflow (Abadi et al. 2016). We use the following methodology to drive an evaluation.



**Figure 9.** A CDF plot that shows the score distribution as a function of the  $PC_{score}$  threshold for the IS, V, J classes on a per-sector basis. As illustrated by the vertical dotted line, Nigraha assigns a low score ( $PC_{score} < 0.4$ ) to nearly 95 per cent of the false positives.

#### 3.3.1 Model training

We train our models using the labeled data sets described in Section 2.4. The labels corresponding to the positive and negative examples used in the training are defined in Table 1. For model training and evaluation, as is customary with training machine learning models, we randomly split the labeled data into three sets: namely, 80 per cent is the *training set* used for model training, 10 per cent is the *validation set* used for model validation, and 10 per cent is set aside as the *test set* which is used solely for model evaluation. Additionally, we augment the test set by injecting TCEs corresponding to known planets and eclipsing binaries (see Section 3.4). We train our models using a batch size of 64 with the Adam optimizer (Kingma & Ba 2014). We enabled early stopping in the training step and the models were trained for roughly 45–60 epochs. We tuned a few of the hyperparameters of the Adam optimizer and found that the best models were with values for learning rate  $\alpha = 1e - 5$  and  $\epsilon = 1e - 8$ .

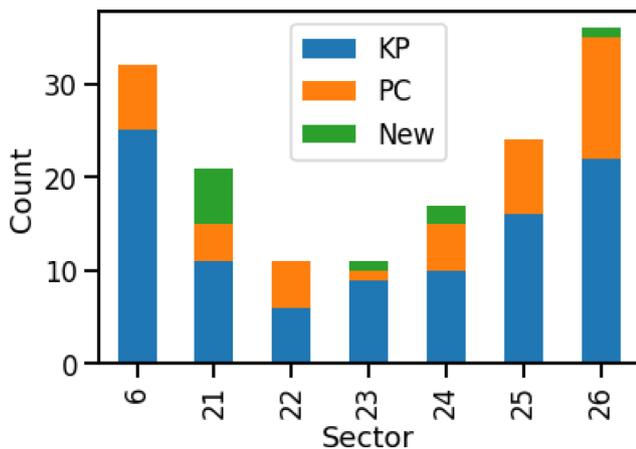
It is well known that model training with imbalanced distributions is hard as the model biases its learning towards features of the negative examples (Chawla 2005). We took two steps to remedy this problem. First, we initialize the output bias in the `sigmoid` output layer to the ratio of the positive to negative examples [i.e.  $\log(pos/neg)$ ]. Secondly, we use balanced-batch sampling which is a well-known oversampling technique to train the model. With balanced-batch sampling, each mini-batch contains an equal number of positive and negative examples (He & Garcia 2009).

#### 3.3.2 Generating predictions

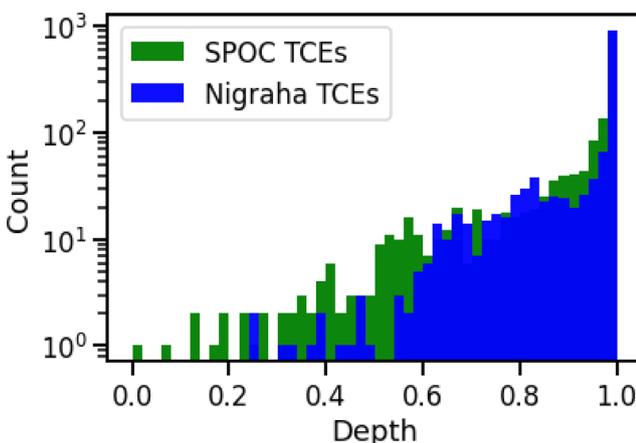
We use ‘model averaging’ to generate predictions. That is, we train 10 independent models on the same training set using random parameter initialization and then compute the likelihood of the TCE being a planet as the average of the score predicted by each model. Model averaging is a well-known ensemble learning technique that is commonly used to minimize variance in a neural network model performance (Karpathy 2019).

**Table 2.** The flow of numbers for Sectors 6 and 21–26, and deduplicated totals through the different stages of our pipeline. (i) Starting from ~20000 targets that all TESS sectors have, the numbers that remain after different steps are as follows: (ii) Science Processing Operations Center/Quick Look pipeline (SPOC/QLP) found TCEs using BLS, (iii) planet candidates and false signals from followed-up SPOC TCEs (PC includes known planets and confirmed planets; FS includes false positives and false alarms), (iv) Nigraha TCE objects after excluding those with missing stellar parameters. These are more numerous than SPOC TCEs for some sectors likely due to TLS catching shallower transits, (v) the number of TCEs common to SPOC and Nigraha, (vi) new Nigraha PCs, (vii) PCs with SPOC overlap (these have SPOC DV reports), (viii) DAVE vetting for the SPOC overlap PCs (FPs are false positives), (ix) Nigraha-only PCs, (x) DAVE vetting for Nigraha-only PC. The last column of Rows 6 and 7 denotes sums from the other columns with and without duplicates.

Row number	Flow step	S6	S21	S22	S23	S24	S25	S26	Total
1	TPS transit searches	19 995	19 995	19 999	19 998	19 992	19 994	19 992	82 647
2	SPOC TCE objects	812	638	561	501	956	974	1012	3724
3	TESS (PC FS)	81 32	101 15	78 14	77 15	116 18	123 16	125 15	358 65
4	Nigraha TCEs	710	771	608	571	928	1009	977	5574
5	common TCE objects	388	340	275	278	604	603	663	3151
6	Nigraha TCE ( $PC_{score} > 0.4$ )	47	42	26	19	60	57	55	306/230
7	Nigraha TCE with SPOC overlap	42	23	21	16	53	51	50	256/181
8	DAVE vetted (PC FP)	7 35	4 19	5 16	1 15	5 48	8 43	13 37	43 213
9	Nigraha-only TCEs	5	19	5	3	7	6	5	50
10	DAVE vetted Nigraha-only (PC FP)	0 5	6 13	0 5	1 2	1 6	0 6	1 4	9 41



**Figure 10.** We identify planet candidates based on the predictions generated by our model. Since this includes the previously known ones, the histogram provides a breakdown for each of the sectors we evaluated. KP are known planets, PC are planet candidates that are also common to SPOC (Tables 3 and 4), and New are vetted Nigraha-only planet candidates (Table 5) found by our ML pipeline.



**Figure 11.** SPOC and Nigraha TCEs. For the SPOC part, we have used their transit depth in ppm - it varies ever so slightly with Nigraha depths. Nigraha TCEs at shallower depths are far more numerous. Here and in the next two figures, depth is described as the fraction of unobscured light.

### 3.4 Evaluation

Our test set consists of 632 TCEs of which 56 TCEs have one of KP, PC, CP labels corresponding to the positive instances (i.e. only 8.8 per cent are positive). To address the imbalance a little, we augment the test set with an additional 175 labeled TCEs from Sector 6 and Sectors 21–26. The labels for these TCEs are from the TESS TOI catalogue and were assigned based on manual vetting by the TESS team (Guerrero et al. 2019). This augmented data set consists of (i) 98 TCEs that are positive instances which correspond to either previously known planets (i.e. KP) or planets that were confirmed through followup of TESS candidates (i.e. CP), and (ii) 77 TCEs that are negative instances which have false positive labels, namely, FP, EB. With this augmentation, the combined test set now has 19.1 per cent positive instances.

With imbalanced data sets, it is generally recommended that the desirable metrics for evaluating model performance are AUC and precision-recall curves (Davis & Goadrich 2006). We evaluate these two metrics in the following subsections. Additionally, we examine how our model scores on the various TCE classes (KP, CP, PC) in the augmented test set. For this analysis, we also include the PC's from the TOI catalogue from Sectors 1–5 and 7–20.

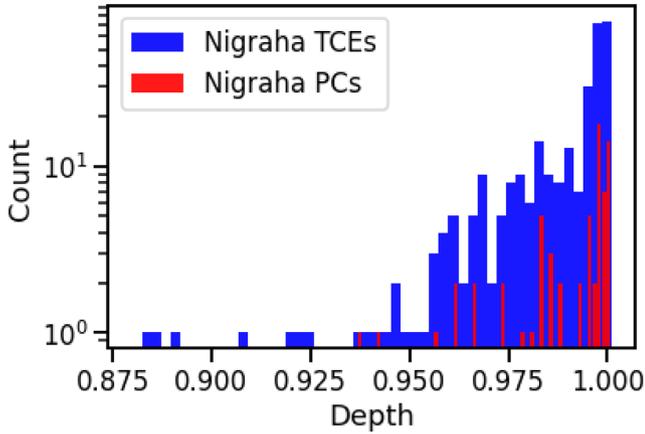
#### 3.4.1 AUC

The AUC metric corresponds to the area under the receiver-operator characteristic curve which is equivalent to the probability that a randomly selected positive is ranked above a randomly selected negative. The AUC score for our model on the combined test set is 0.908.

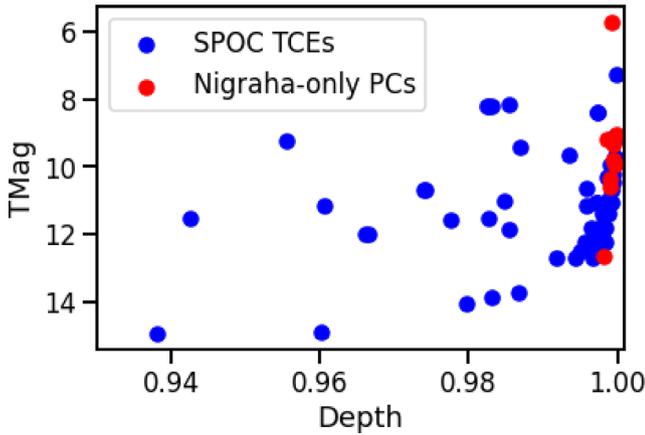
#### 3.4.2 Setting $PC_{score}$ threshold

Fig. 5 shows the precision-recall curves as a function of threshold ( $PC_{score}$ ) for the test set. For the analysis in this section, we choose the threshold by picking a point closest to the intersection of the two curves as it represents a trade-off between precision and recall. As the figure shows, this point corresponds to a threshold of  $PC_{score} > 0.7$  and we use this for candidate selection.

We next construct a confusion matrix for the test set at this threshold (see Fig. 6). In the figure, we use PC to denote TCEs corresponding to the positive instances and FP to denote the negative



**Figure 12.** Distribution of Nigraha TCEs (blue) which get filtered down to candidates (PC's, shown in red). For completeness depths from multiple sectors are included when same objects were found identified multiple times.



**Figure 13.** Nigraha-only PCs (red) have shallower transits than SPOC TCEs (blue).

instances. The model has an accuracy of 87.2 per cent (fraction of the instances that are correctly classified), a precision of 88.8 per cent (fraction of the positive instances that are indeed true positives), and a recall of 74.3 per cent (fraction of the true positives that are recovered).

### 3.4.3 Model scoring

A question of interest is the  $PC_{\text{score}}$  range our model assigns to TCEs corresponding to various classes. For this analysis, we also include previously classified PC's from the TOI catalogue. Fig. 7 shows the cumulative distribution function (CDF) of the fraction of TCEs recovered above a given  $PC_{\text{score}}$ . Our model recovers 91 of 99 (91.9 per cent) of the KP's and CP's with  $PC_{\text{score}} > 0.7$ . Such a high recall at a high threshold is desirable. However, for the same threshold, over 30 per cent of the false positives also get included. Such misclassified false positives can be eliminated by downstream vetting.

Given the threshold of 0.7, we find that our model fails to recover eight TCEs. Of these, six are KP's and one is a CP for which our model predicts  $PC_{\text{score}}$  in the range of 0.44–0.68. We visually examined the light curves for the missing objects and did not find any systematic issues. It is reasonable to expect that as the quality of training data

improves these objects would be scored higher. A CP we fail to recover (TIC 260128333) is assigned a score of  $PC_{\text{score}} = 0.039$ . For this object, the Group Comments column in the TESS TOI catalogue observes that ‘some transits may be spurious so period could be at longer interval’.

Next, Fig. 8 shows the CDF of the scores for the various classes of false positives in the test set. On this figure, we also overlay the scores assigned to the PC's from the test set. Note that, the false positives are in general assigned much lower scores by Nigraha. In particular, note that, for the IS, V, J classes, over 90 per cent of those TCEs are scored at 0.2 or less.

Finally, as the noise properties and scattered light vary on per-sector basis, it is possible that J, IS may be sector dependent. This is of interest as it could potentially affect our results: we train our model with data from Sectors 1–5 and then use that to predict on another sectors. To understand this effect, Fig. 9 shows the CDF of the scores for the IS, V, J classes (from our test set) on a per-sector basis. As the graph shows, we find that across these sectors scores do not exhibit substantial variation. That is, we find that nearly 95 per cent of the FP's in each sector are scored 0.4 or lower. This is desirable as it suggests that Nigraha is able to consistently rule out vast majority of the false positives.

To summarize, picking the threshold is a precision-recall trade-off: it can be moved in either direction to increase or decrease recall or precision at the cost of the other depending on one's aim. Thus, we could lower the threshold to 0.4, and recover the missing KP's, but at the expense of increasing false positives. The false positives then will have to be excluded by the downstream vetting.

### 3.5 Comparison to prior models for TESS

We find that Nigraha recovers 19 of the 22 TCEs with KP label from Sector 6 with  $PC_{\text{score}} > 0.8$ . In contrast, Astronet-Vetting recovers seven of these (see table 2 of Yu et al. 2019). Osborn et al. (2020) trained on simulated data and applied their model to TESS Sectors 1–4, whereas our training set included data from TESS Sectors 1–4. Consequently, we can not do a direct comparison; but their recall is 61 per cent compared to Nigraha's 74.3 per cent. These non-exhaustive comparisons suggest that Nigraha generalizes better.

## 4 GENERATING NEW CANDIDATES

We use the *inference* path of our pipeline (see Fig. 3) to identify candidates from Sector 6 and Sectors 21–26. We begin by providing an overview how the various steps in the pipeline prune the candidates (Table 2) and then describe our results.

### 4.1 Overview

Our TCEs comprise the ones that have been found by the TESS Science Processing Operations Center (SPOC) pipeline (Jenkins et al. 2016) as well as previously unseen TCEs. The new TCEs we identify are due to the difference in the search procedure: While we use TLS and it is better at recovering shallow transits, the SPOC pipeline uses a wavelet-based adaptive matching filter for the search (Jenkins et al. 2010) and may have missed recovering those TCEs. The effect of this difference is illustrated in Table 2 (Rows 2, 4, and 5). The inputs to both the pipelines is the same: Row 1 in the table shows the number of unique light curves for the TESS targets across the various sectors. From the light curves of these targets, Row 2 and Row 4, respectively, show the number of TCEs found by SPOC and Nigraha. In anticipation of the inputs needed by our ML model, we

**Table 3.** PCs from TCEs in Sectors 6 and 21–26 that overlap the SPOC TCEs. These are DAVE and Gaia vetted. Various other metadata are provided to allow sub-selection for follow-up observing.  $T_0 = \text{BJD} - 2457000$ . Note that we do not account for light contamination from nearby stars when estimating planet radii. Hence, our computed value is likely to be higher. We take a conservative approach to flag all candidates whose planet radius exceeds 13 by labeling them in red.

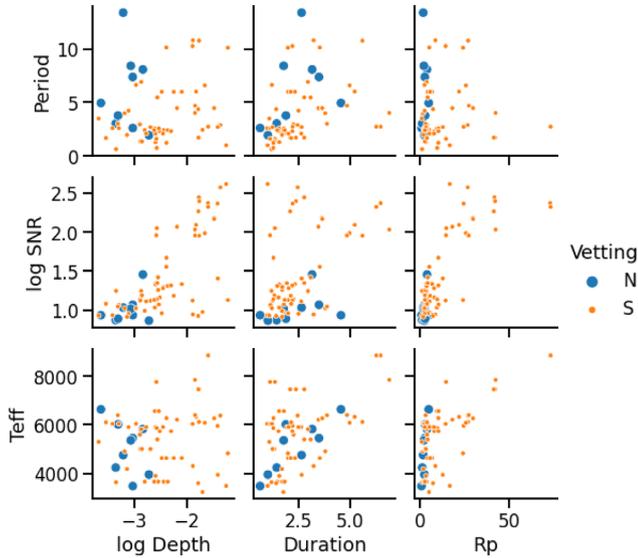
TIC/sector	RA	DEC	TMag	Depth (ppm)	Epoch	Duration (h)	Period (d)	$R_p$ ( $R_\oplus$ )	$R_\star$ ( $R_\star$ )	$PC_{\text{score}}$	RUWE
167554898/6	101.3158	−61.0915	9.663	6579	1470.931	3.468	$4.452 \pm 0.025$	18.320	2.068	0.811	1.29
123201406/26	280.9439	42.8392	11.165	4137	2010.889	1.262	$1.203 \pm 0.004$	10.478	1.492	0.609	1.34
279569731/6	105.9310	−57.6608	10.442	654	1471.519	3.806	$3.980 \pm 0.033$	1.975	0.707	0.538	0.87

**Table 4.** PCs from TCEs in Sectors 6 and 21–26 that overlap the SPOC TCEs. Though these pass vetting by DAVE, they however fail some of the diagnostic tests that were run by the SPOC data validation pipeline. Since those tests are known to produce false negatives, we list these PC's separately as many of these may still be viable for follow-up. Also, we targets with a large radius ( $>13$ ), short period ( $<1$ day), and high RUWE score ( $>1.4$ ) are labeled in red. Various other metadata are provided to allow sub-selection for follow-up observing.  $T_0 = \text{BJD} - 2457000$ .

TIC/Sector	RA	DEC	TMag	Depth (ppm)	Epoch	Duration (h)	Period (d)	$R_p$ ( $R_\oplus$ )	$R_\star$ ( $R_\star$ )	$PC_{\text{score}}$	RUWE
388076435/25	317.8838	68.4115	9.176	509	1984.819	1.117	$2.576 \pm 0.007$	2.802	1.137	0.996	1.08
374860586/6	88.6250	−62.6348	10.974	1279	1470.574	2.247	$2.933 \pm 0.017$	4.181	1.070	0.993	0.84
298969838/26	288.5107	51.1625	14.965	61944	2019.446	1.970	$10.115 \pm 0.056$	24.324	0.895	0.986	1.14
264679887/6	82.1150	1.3003	10.305	1356	1468.840	2.956	$6.914 \pm 0.038$	3.956	0.984	0.971	1.16
275633324/23	294.0727	57.2742	11.809	2563	1929.949	1.549	$1.266 \pm 0.005$	3.527	0.638	0.895	5.45
199662939/24	251.4267	54.5447	12.510	3075	1956.295	1.507	$2.293 \pm 0.010$	8.197	1.354	0.895	1.09
272781944/26	238.1813	68.7650	9.085	287	2011.568	2.131	$1.652 \pm 0.010$	4.146	2.240	0.778	0.96
256324377/21	297.8002	62.2866	9.688	448	1870.822	1.167	$0.611 \pm 0.002$	1.350	0.584	0.730	0.93
235125815/6	102.1588	10.1983	14.053	20262	1469.040	1.749	$2.423 \pm 0.013$	5.252	0.338	0.702	4.63
46627823/6	85.1816	−17.5457	11.002	15258	1470.504	5.247	$6.011 \pm 0.044$	22.091	1.638	0.700	1.12
233195052/25	264.2525	67.0935	14.887	39835	1985.136	2.647	$2.679 \pm 0.020$	12.957	0.594	0.668	0.99
143114211/25	257.0213	17.9605	9.924	1082	1988.147	1.185	$6.636 \pm 0.018$	2.696	0.750	0.662	4.07
164458426/26	283.1663	45.4030	11.872	14640	2010.796	4.868	$4.695 \pm 0.036$	27.560	2.086	0.647	0.97
138819130/22	175.4461	25.6577	11.508	17369	1909.557	5.623	$10.781 \pm 0.113$	27.187	1.889	0.599	0.91
258177087/24	239.6334	27.7400	7.290	208	1956.702	1.060	$3.484 \pm 0.015$	1.212	0.769	0.564	1.03
236788376/26	283.2101	56.0622	13.729	13243	2019.902	3.232	$10.810 \pm 0.085$	8.810	0.701	0.555	1.00
165453878/22	200.8407	50.8268	11.165	39322	1902.021	2.202	$2.706 \pm 0.005$	24.694	1.140	0.546	0.97
117280404/21	131.1200	28.4564	9.404	13059	1876.781	2.190	$10.274 \pm 0.041$	14.558	1.167	0.514	1.02
259006174/26	290.3177	70.2361	11.057	725	2010.483	1.684	$1.939 \pm 0.007$	2.064	0.702	0.457	0.79
232612421/26	259.0272	63.1145	13.879	16792	2014.134	1.396	$4.361 \pm 0.018$	7.499	0.530	0.447	1.05
167696018/6	102.8952	−62.7181	10.649	4122	1468.363	3.537	$10.140 \pm 0.115$	5.426	0.774	0.441	0.88
368260545/26	265.4772	21.7358	11.587	22472	2015.478	4.997	$6.578 \pm 0.050$	25.965	1.586	0.437	1.15
468739039/25	20.3185	74.0719	10.208	485	1985.549	3.886	$2.992 \pm 0.023$	3.290	1.367	0.423	1.00
235726293/26	276.4707	24.4778	11.508	57327	2011.717	0.975	$0.977 \pm 0.002$	16.929	0.647	0.417	1.67
461634737/26	18.0993	88.3591	10.706	778	2011.987	2.491	$4.220 \pm 0.031$	2.774	0.910	0.402	1.10
233051157/26	268.7985	61.1053	12.190	2693	2010.859	1.256	$0.720 \pm 0.003$	9.831	1.735	0.401	1.01

**Table 5.** As in Table 3 but for Nigraha-only PCs. These are vetted with DAVE.  $T_0 = \text{BJD} - 2457000$ . TIC 392365181 (Sector 21) has a companion 5 arcsec away. It is over 3.7 mag fainter and with a RUWE score of 1.07, and unlikely to have contributed significantly to the variability of TIC 392365181. TIC 468922064 (Sector 26) has a companion that 3.5 mag fainter with high RUWE a few arcsec away. The nearest Gaia object to TIC 1400994843 is over 7 arcsec away and has a RUWE of 1.05.

TIC	RA	DEC	TMag	Depth (ppm)	Epoch ( $T_0$ )	Duration (h)	Period (d)	$R_p$ ( $R_\oplus$ )	$R_\star$ ( $R_\star$ )	$PC_{\text{score}}$	RUWE
407591297/21	152.1755	35.5470	10.574	936	1870.655	0.607	$2.594 \pm 0.007$	1.046	0.313	0.859	1.16
392365181/21	155.9795	44.1149	9.954	441	1871.029	1.420	$3.014 \pm 0.012$	1.461	0.637	0.655	0.90
199608547/23	248.7099	56.7093	12.656	1909	1931.255	0.988	$1.917 \pm 0.007$	2.680	0.562	0.628	1.08
424731749/21	255.1672	60.4345	9.183	1461	1871.163	3.159	$8.078 \pm 0.037$	4.114	0.986	0.611	1.08
166012100/21	212.0799	64.0684	9.313	613	1870.525	2.640	$13.405 \pm 0.044$	1.995	0.738	0.600	0.95
199610508/21	249.1159	58.5853	10.372	934	1876.602	3.495	$7.375 \pm 0.079$	2.889	0.866	0.568	1.29
468922064/26	96.6825	70.8227	9.779	488	2013.372	1.872	$3.765 \pm 0.015$	3.031	1.255	0.553	1.11
445810510/21	170.4891	50.6882	9.067	229	1871.725	4.574	$4.940 \pm 0.035$	5.092	3.077	0.457	2.69
1400994843/24	261.2504	67.3067	5.717	858	1958.469	1.761	$8.421 \pm 0.049$	2.378	0.743	0.401	*



**Figure 14.** Nigraha-only PCs (N) and those overlapping with SPOC (S).

**Table 6.** Comparison of SNR/SDE for the Nigraha-only TCEs between BLS and TLS.

TIC	Sector	BLS		TLS	
		SNR	SDE	SNR	SDE
407591297	21	7.487	13.429	8.610	17.853
392365181	21	9.223	9.672	7.458	13.388
199608547	23	4.525	9.387	7.344	9.923
424731749	21	26.640	9.355	28.623	10.167
166012100	21	7.036	4.723	10.774	9.590
199610508	21	13.252	6.296	11.700	9.149
468922064	26	6.449	10.912	7.775	10.406
445810510	21	7.552	7.810	8.606	11.531
1400994843	24	35.305	8.983	10.424	11.368

exclude objects that are missing at least one of the stellar parameters (described in Section 2.5). Row 5 shows the TCEs that are common to both the pipelines.

We use our trained model to generate a prediction (i.e.  $PC_{\text{score}}$ ) for each of the Nigraha found TCEs. In this section, we apply a threshold ( $PC_{\text{score}} > 0.4$ ) to identify a promising sample of candidates for further vetting. We use Fig. 7 to pick this threshold for which there is a high recall (at least 90 per cent), and we use downstream vetting to eliminate false positives. The candidate set we construct includes objects that have been previously identified either by the TESS pipelines or by others in the community. Row 3 of Table 2 shows the number of labeled objects across the various sectors that we found by cross-referencing with the publicly available TESS TOI and TESS CTOI catalogues.<sup>8</sup> Next, any candidates that are left after we filter out these known ones are *new*. Row 6 shows the count of the *new* planetary candidates. Lastly, we perform preliminary vetting (Section 4.3) using a series of steps including evaluation with DAVE (Discovery and Validation of Exoplanets; Kostov et al. 2019a) and cross-checking with GAIA to rank our candidates. The number of candidates subject to vetting is shown in Rows 7–10 of Table 2. The final confirmation of these candidates as bona fide planets is only

<sup>8</sup><https://exofop.ipac.caltech.edu/tess/> accessed 10/12/20

through specialized observations. Hence, our aim with vetting is to minimize the number of candidates to the most promising ones.

## 4.2 Finding candidates

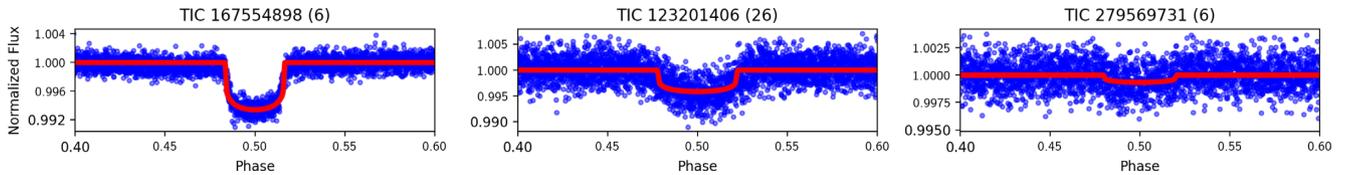
Fig. 10 shows the number of TCEs in each of the sectors with a  $PC_{\text{score}} > 0.4$  as predicted by our model. The figure also provides a breakdown of the previously known planets (KP + CP) and candidates (PC) from the TESS TOI/CTOI catalogues. After accounting for the known objects, we are left with a candidate sample that is roughly between 10 and 40 new candidates per sector (Row 6 of Table 2). Fig. 11 shows the distribution of TCEs as a function of transit depth. As expected, we find that Nigraha-only TCEs with shallow transit depths to be numerous. Next, when apply our threshold for candidate selection ( $PC_{\text{score}} > 0.4$ ), the number of candidates reduces drastically. This is illustrated in Fig. 12. This candidate sample includes TCEs found by both SPOC pipeline as well as Nigraha-only. Fig. 13 shows that Nigraha-only PC's tend to be fainter and are likely to have shallower transits compared to SPOC TCEs. In what follows, we perform preliminary vetting of these candidates and then describe their characteristics.

## 4.3 Preliminary vetting of new candidates

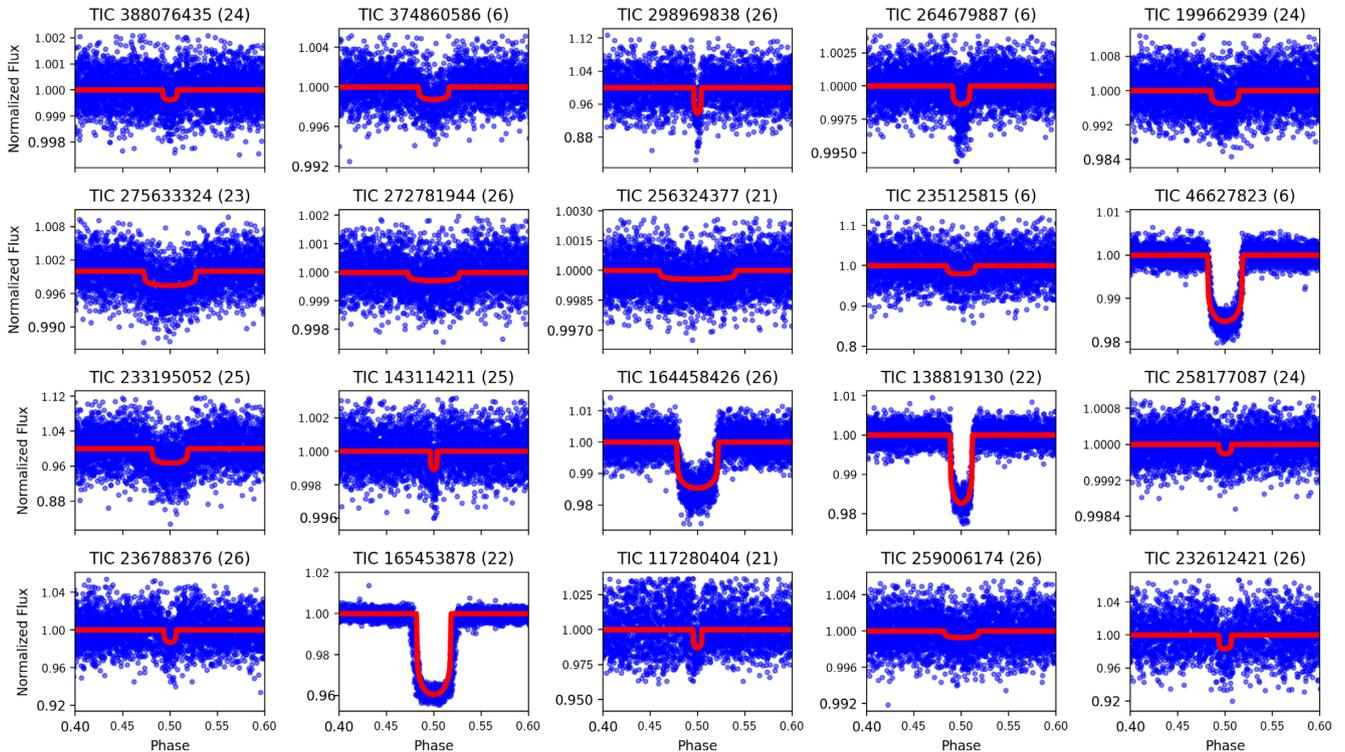
The next step in validating planet candidates as bona fide planets is to either perform follow-up observations using telescope or use statistical methods. Since telescope time is scarce, statistical validation is becoming an increasingly viable alternative. For instance, with Kepler/K2, *vespa* (Morton 2012; Morton et al. 2016) has been extensively used to authoritatively validate hundreds of candidates as *bona fide* planets. Such an authoritative tool for TESS is not yet available. However, it is possible to do some preliminary vetting using DAVE (Kostov et al. 2019a). DAVE, which was originally developed for K2, has been adapted and extensively used for vetting TESS candidates (Kostov et al. 2019b; Gilbert et al. 2020).

Briefly, DAVE can be configured to perform a set of tests to determine if the transit-like signals are due to common false positive scenarios; using the data validation light curve, it checks for odd-even differences between consecutive transits, secondary eclipses, stellar variability mimicking a transit; using the target pixel file, it checks for photocenter shifts during transit. Both of these inputs to DAVE are TESS products which are available for public download from MAST. Based on these tests, DAVE provides a preliminary disposition: a TCE is either identified as a candidate or it is flagged as a false positive with a diagnostic message. For example, Fig. 17 shows the summary report for TIC 167554898 that DAVE identified as a candidate.

The data validation (DV) products (namely, light curves and reports) are available only for the TCEs identified by the SPOC pipeline. This is an issue for us since our method is capable of identifying TCEs not in the SPOC set (such as, those with shallow transits). Hence, of the 243 candidates we identify, we were able to vet 211 TCEs across these sectors (line 7, Table 2); of these, DAVE provided a positive disposition of *candidate* for 21 TCEs (line 8, Table 2). The false positives consisted of: 106 TCEs that failed the ModelShift test, 45 TCEs that failed the odd-even depth test, and 63 TCEs that had centroid offset changes in the photometric image. Additionally, there are instances where a TCE was classified by DAVE as a candidate in one sector, but flagged as a false positive in another sector. For such cases, we consider the TCE as a false positive. Conversely, there are instances where a TCE was classified



**Figure 15.** The binned/phase folded light curves for the new planet candidates found by our model and subsequently vetted with DAVE. The SPOC TCEs in this panel pass all vetting tests and are not flagged in the DV report.



**Figure 16.** The binned/phase folded light curves along with the transit model (shown in red) for the top 20 new planet candidates found by our model and subsequently vetted with DAVE. *The SPOC TCEs in this panel have been flagged in the DV report.*

by Nigraha as candidate in one sector but not in another. For such cases, we do not consider the TCE as part of the candidate sample.

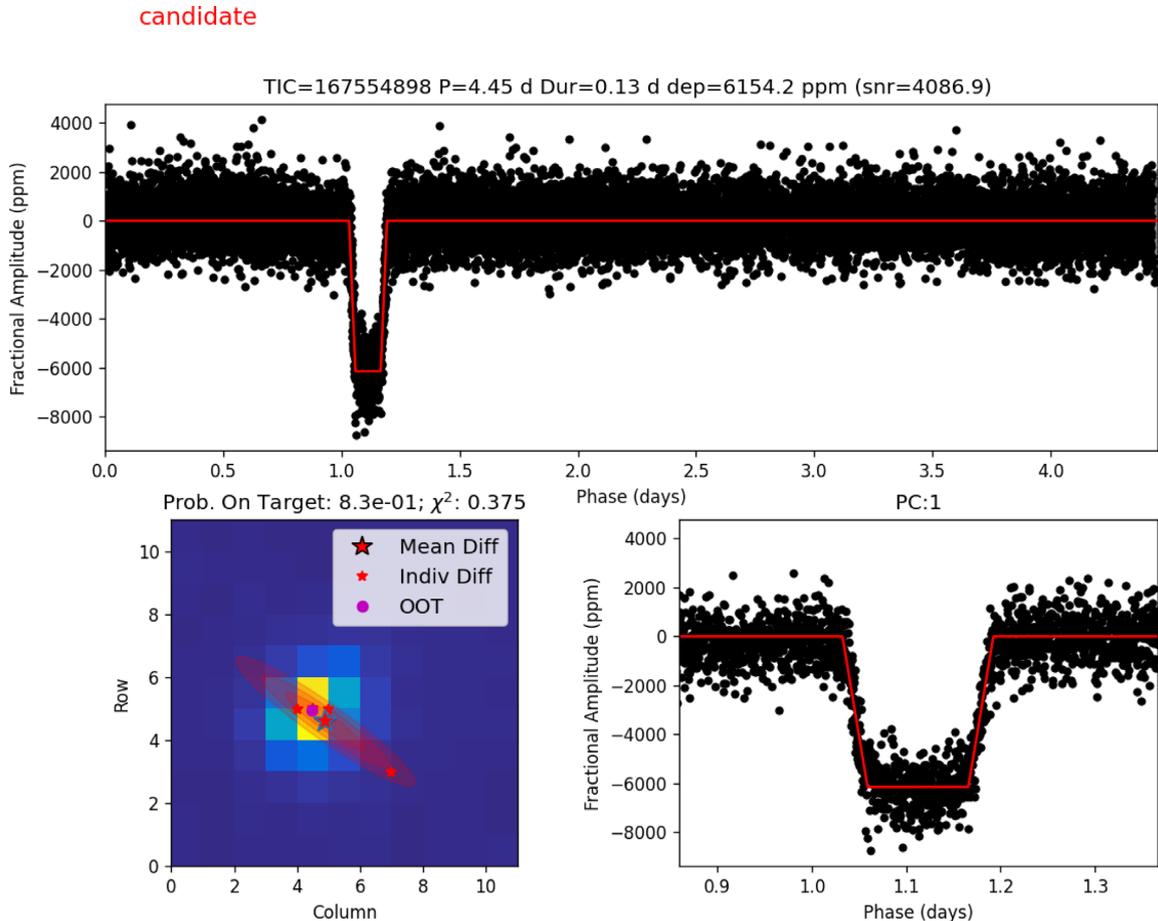
For these SPOC TCEs that are identified as a planet candidate by DAVE, we perform an additional vetting step. For each target, we cross-check with the associated SPOC generated data validation report as well as the TESS Exoclass catalogue (Burke et al. 2020) to see if the diagnostic tests have flagged any issues such as, centroid shift, planet radius is too big, etc. Table 3 lists the set of TCEs that pass both the vetting steps. The diagnostic tests run as part of the SPOC data validation steps are known to produce false negatives. For instance, the data validation report for TIC 103633434 flags a centroid out-of-transit test, but the TCE corresponds to a CP. Similarly, the TCEs from several other bona fide planets in TOI catalogue have been flagged as part of the data validation checks. Hence, we report those TCEs that did not pass the additional vetting separately and they are listed in Table 4.

We vet the remaining Nigraha-only candidates (line 9, Table 2) with DAVE by making a few modifications. Since the DV products are not available for these targets, we input the PDCSAP\_FLUX light curves after detrending (as described in Section 2) along with our transit parameters to DAVE. Of the 51 Nigraha-only TCEs, we find that 3 pass all the tests and the remaining 48 are false positives. These

false positives consisted of: 37 TCEs which failed the ModelShift test, 10 TCEs which failed the odd-even depth test, and 1 TCE which had centroid offset changes in the photometric image. Table 5 lists the set of Nigraha TCEs that are identified as candidates by DAVE.

Figs 15, 16, and 18, respectively, show the phase-folded light curves for all of the candidates from Tables 3, 4, and 5 which were identified by our pipeline. We further checked in the Gaia catalogue the RUWE associated with each source. The RUWE is expected to be around 1.0 for single stars, and a value significantly greater, e.g.  $> 1.4$ , possible binarity or an issue with the astrometric solution (Lindgren et al. 2018).

Fig. 14 shows some of the planetary characteristics for our vetted/unvetted candidate sample set. In this figure, we compute a planet's radius as an approximation using the relation,  $\delta_{\text{tra}} \approx k^2$ , where  $\delta_{\text{tra}}$  is the loss of light during transit and  $k = \frac{R_p}{R_*}$ . Note that by using this approximation, we neither account for light contamination from nearby stars nor light from the planetary nightside that is typically negligible (Winn 2014). From Row 1 and 3 of this panel, we find that the final set of candidates we are left with are super-Earths ( $R_p < 4R_{\oplus}$ ) with short period and duration, suggesting close proximity of their orbit to the host star.



**Figure 17.** Example summary report of vetting by DAVE. This is for our candidate TIC 167554898.

#### 4.4 Discussion: BLS versus TLS

Some of the new candidates we identified above are from light curves that were not identified as containing TCEs by the SPOC pipeline. We believe the underlying reason that the TCEs were missed is primarily due to the low SNR/SDE in the transit search with BLS. Note that the widely used thresholds for BLS are  $\text{SNR}_{\text{BLS}} > 7.1$ ,  $\text{SDE}_{\text{BLS}} > 9$  (Crossfield et al. 2016; Vanderburg et al. 2016). Table 6 shows the results of our analysis. We find that six of the nine Nigraha-only TCEs were not recovered by BLS.

Interestingly, we find that TIC 388076435 is recovered as a candidate by Nigrahain Sector 24, whereas SPOC does not. In this sector, our TLS search returned a depth of 383.76 ppm at a SNR of 8.6 and SDE of 11.8. However, we find that the same TIC is also identified as a TCE by SPOC in Sector 25 (see Table 4) with a depth of  $562.4 \pm 66.5$  ppm (as per the DV report), while our TLS search returned 509 ppm.

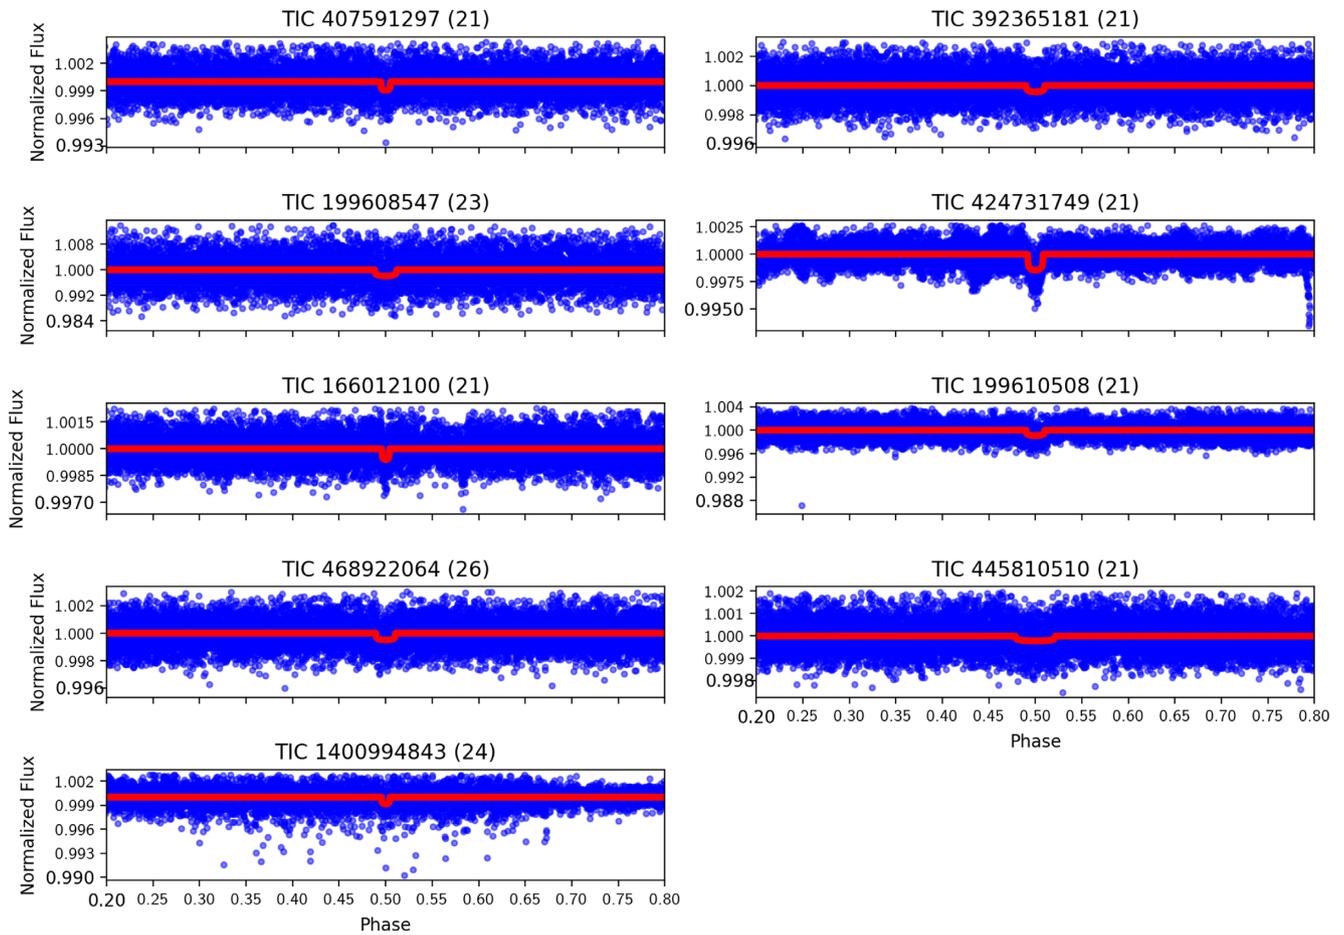
## 5 CONCLUDING REMARKS

In large-scale sky surveys for detecting exoplanets, it is well known that the number of bona fide planetary candidates is dwarfed by the number of false positives in the observed data. For instance, with TESS, the expected yield is about 80 candidates per TESS sector<sup>9</sup> that comprises 20000 targets. Using machine learning techniques

to build models that enable rapid classification has proven to be a promising approach (Ansdell et al. 2018; Shallue & Vanderburg 2018; Schanche et al. 2019; Yu et al. 2019; Osborn et al. 2020). The TESS team operates multiple pipelines (namely, SPOC and QLP) to identify promising candidates, several of which have been validated as bona fide planets.

As a complementary effort, we set out to build a machine-learning-based pipeline, *Nigraha*, using novel approaches from ground-up. First, the method used to search for transit signals in a light curves (TLS) differs from prior approaches (such as MIT QLP for TESS) which use BLS and its variants. Compared to BLS, TLS has better sensitivity towards recovering shallow transits caused by near-Earth-sized planets. Consequently, this expands the search space for identifying candidates. Secondly, to enable the model to better distinguish planetary candidates from false positives caused by eclipsing binaries, we use a novel input representation of the light curve. We trained our model using data from TESS sectors and evaluated it. *Nigraha* has an AUC score of 0.908. In a test comprising light curves from known and TESS confirmed planets, we find that our model recovers 91 of 99 known/TESS confirmed planets. We then used our trained model to generate predictions for data from various TESS sectors. Besides recovering known planets and planet candidates, our model identified new candidates. We performed preliminary vetting of all candidates using a sequence of steps including DAVE to rule out false positive scenarios and checked for binarity using the RUWE parameter from GAIA. Although neither of these methods are foolproof, by minimizing the number of candidates

<sup>9</sup><https://tess.mit.edu/toi-releases/toi-release-faqs/>



**Figure 18.** The binned/phase folded light curves along with the transit model (shown in red) for the new planet candidates found by our model. The corresponding TCEs were not in the SPOC list.

we make them better candidates for follow-up. This is true for the shallow candidates we discovered as well as the existing planet candidates we independently affirm. The value of finding vetted shallower transits is demonstrated by TIC 388076435. Nigraha finds it in Sector 24, and only using data in Sector 25 does it become a SPOC candidate. Our Tables 3, 4, and 5 summarize the best cases in Sectors 6 and 21–26 including a few with reservations.

Some of these may yet turn out to be false positives (possibly diluted eclipsing binaries), but a subset of those will likely be binaries with very faint companions, perhaps even white dwarfs, and thus interesting in their own right. In the near future we will be extending our model to support multiclass classification with eclipsing binaries being an explicit class, and variables another.

Our procedure for finding TCEs has been intentionally conservative, and there are a few ways to strengthen our approach to increase the candidate yield. For instance, we only search for a single TCE per light curve; we process light curves on a per-sector basis; we discard candidates for which the stellar/transit parameters are not available. We intend to address these issues as part of ongoing work in this project. Furthermore, this should also help with identifying both stars with multiplanet systems as well as planets with longer periods (such as, those with periods greater than 27 d and discovered by combining data from multiple sectors). We are also exploring ways to improve model performance based on ideas, such as model training using data from Sector 6 and later as that data has improved camera pointing systematics. There are also alternate

input data representation and machine learning model architecture optimizations to be tried.

We have released the code and models developed for this paper as open source.<sup>10</sup> We intend to submit the candidates found by Nigraha pipeline to the TESS Exoplanet Followup Program (ExoFOP-TESS). Independently, we are proposing high-resolution, time-resolved photometry for the most promising of our candidates using WIRC with the Engineered Diffuser (Stefansson et al. 2017; Vissapragada et al. 2020) and dual-band CHIMERA (Harding et al. 2016) at the Palomar 5 m telescope to rule out false positives and confirm the targets as planets. We hope other teams will also observe these candidates.

## ACKNOWLEDGEMENTS

We thank the anonymous referee whose detailed feedback greatly improved this paper. We thank David Ciardi, Joe Ninan Philip, and Shreyas Vissapragada for insightful comments and feedback on various aspects of this work. We also thank Tim Morton for discussions and feedback on early drafts of this paper.

Funding for the TESS mission is provided by NASA’s Science Mission directorate. We acknowledge the use of public TESS Alert data from pipelines at the TESS Science Office and at the TESS Science Processing Operations Center. This research has made use

<sup>10</sup><https://github.com/ExoplanetML/Nigraha>

of the Exoplanet Follow-up Observation Program website, which is operated by the California Institute of Technology, under contract with the National Aeronautics and Space Administration under the Exoplanet Exploration Program. This paper includes data collected by the TESS mission, which are publicly available from the Mikulski Archive for Space Telescopes (MAST).

This research made use of LIGHTKURVE, a Python package for Kepler and TESS data analysis (Lightkurve Collaboration et al. 2018).

This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement.

AM acknowledges support from the National Science Foundation (NSF Grant 1640818).

## DATA AVAILABILITY

The data underlying this article is available for public download on the GitHub repository at <https://github.com/ExoplanetML/Nigraha>. Additionally, on the same GitHub repository, we have released the code and models developed as part of this work.

## REFERENCES

- Abadi M. et al., 2016, preprint ([arXiv:1605.08695](https://arxiv.org/abs/1605.08695))
- Ansdell M. et al., 2018, *ApJ*, 869, L7
- Armstrong D. J., Pollacco D., Santerne A., 2016, *MNRAS*, 465, 2634
- Armstrong D. J. et al., 2018, *MNRAS*, 478, 4225
- Burke C. et al., 2020, TESS Exoclass. Available at: <https://github.com/christopherburke/TESS-ExoClass>
- Chawla N. V., 2005, *The Data Mining and Knowledge Discovery Handbook*. Springer, USA
- Chollet F. et al., 2015, *Keras*. Available at: <https://keras.io>
- Claret A., 2018, *A&A*, 618, A20
- Crossfield I. J. M. et al., 2016, *ApJS*, 226, 7
- Dattilo A. et al., 2019, *AJ*, 157, 169
- Davis J., Goadrich M., 2006, *Proceedings of the 23rd International Conference on Machine Learning. ICML '06*. Association for Computing Machinery, New York, p. 233
- Gilbert E. A. et al., 2020, *AJ*, 160, 116
- Guerrero N. et al., 2019, TESS Objects of Interest Catalog. Available at: [https://tsc.mit.edu/talks/Natalia\\_Guerrero.pdf](https://tsc.mit.edu/talks/Natalia_Guerrero.pdf)
- Harding L. K. et al., 2016, *MNRAS*, 457, 3036
- He H., Garcia E., 2009, *IEEE Transactions on Knowledge and Data Engineering*. IEEE, USA
- Heller R., Rodenbeck K., Hippke M., 2019a, *A&A*, 625, A31
- Heller R., Hippke M., Rodenbeck K., 2019b, *A&A*, 627, A66
- Heller R., Hippke M., Freudenthal J., Rodenbeck K., Batalha N. M., Bryson S., 2020, *A&A*, 638, A10
- Hippke M., Heller R., 2019, *A&A*, 623, A39
- Hippke M., David T. J., Mulders G. D., Heller R., 2019, *AJ*, 158, 143
- Jenkins J. M. et al., 2010, *ApJ*, 713, L87
- Jenkins J. M. et al., 2016, *Proc. SPIE*, 9913, 99133E
- Karpathy A., 2019, *A Recipe for Training Neural Networks*. Available at: <https://karpathy.github.io/2019/04/25/recipe/>
- Kingma D. P., Ba J., 2014, preprint ([arXiv:1412.6980](https://arxiv.org/abs/1412.6980))
- Kostov V. B. et al., 2019a, *AJ*, 157, 124
- Kostov V. B. et al., 2019b, *AJ*, 158, 32
- Kovács G., Zucker S., Mazeh T., 2002, *A&A*, 391, 369
- Lightkurve Collaboration, 2018, *Astrophysics Source Code Library*, record ascl:1812.013
- Lindgren L. et al., 2018, *A&A*, 616, A2
- Mandel K., Agol E., 2002, *ApJ*, 580, L171
- McCauliff S. D. et al., 2015, *ApJ*, 806, 6
- Morton T. D., 2012, *ApJ*, 761, 6
- Morton T. D., Bryson S. T., Coughlin J. L., Rowe J. F., Ravichandran G., Petigura E. A., Haas M. R., Batalha N. M., 2016, *ApJ*, 822, 86
- Osborn H. P. et al., 2020, *A&A*, 633, A53
- Pearson K. A., Palafox L., Griffith C. A., 2017, *MNRAS*, 474, 478
- Pont F., Zucker S., Queloz D., 2006, *MNRAS*, 373, 231
- Ricker G. R. et al., 2014, *Proc. SPIE*, 9143, 914320
- Schanche N. et al., 2019, *MNRAS*, 483, 5534
- Shallue C. J., Vanderburg A., 2018, *AJ*, 155, 94
- Stassun K. G. et al., 2018, *AJ*, 156, 102
- Stefansson G. et al., 2017, *ApJ*, 848, 9
- Thompson S. E., Mullally F., Coughlin J., Christiansen J. L., Henze C. E., Haas M. R., Burke C. J., 2015, *ApJ*, 812, 46
- Vanderburg A. et al., 2016, *ApJS*, 222, 14
- Vissapragada S. et al., 2020, *AJ*, 159, 108
- Winn J. N., 2014, preprint ([arXiv:1001.2010](https://arxiv.org/abs/1001.2010))
- Yu L. et al., 2019, *AJ*, 158, 25
- Zucker S., Giryes R., 2018, *AJ*, 155, 147

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